arXiv:1511.00778v1 [cs.CC] 3 Nov 2015

Basis Collapse for Holographic Algorithms Over All Domain Sizes

Sitan Chen*
sitanchen@college.harvard.edu

November 4, 2015

Abstract

The theory of holographic algorithms introduced by Valiant represents a novel approach to achieving polynomial-time algorithms for seemingly intractable counting problems via a reduction to counting planar perfect matchings and a linear change of basis. Two fundamental parameters in holographic algorithms are the *domain size* and the *basis size*. Roughly, the domain size is the range of colors involved in the counting problem at hand (e.g. counting graph k-colorings is a problem over domain size k), while the basis size ℓ captures the dimensionality of the representation of those colors. A major open problem has been: for a given k, what is the smallest ℓ for which any holographic algorithm for a problem over domain size k "collapses to" (can be simulated by) a holographic algorithm with basis size ℓ? Cai and Lu showed in 2008 that over domain size 2, basis size 1 suffices, opening the door to an extensive line of work on the structural theory of holographic algorithms over the Boolean domain. Cai and Fu later showed for signatures of full rank that over domain sizes 3 and 4, basis sizes 1 and 2, respectively, suffice, and they conjectured that over domain size k there is a collapse to basis size $\lfloor \log_2 k \rfloor$. In this work, we resolve this conjecture in the affirmative for signatures of full rank for all k.

1 Introduction

1.1 Matchgates and Holographic Algorithms

In [19, 20], Valiant introduced the notion of matchgate computation as a method for classically simulating certain quantum gates in polynomial time and asked whether matchgates could be used to derive other polynomial-time algorithms. Given a counting problem, one might hope that its local combinatorial constraints can be encoded by graph fragments G_i so that there exists a bijection between the number of solutions to that counting problem and the number of perfect matchings of an amalgamation Ω of those graph fragments. The celebrated Fisher-Kasteleyn-Temperley algorithm [13, 17] then computes the latter in polynomial time if Ω is planar.

The drawback of the above approach is that if one seeks a one-to-one reduction to counting perfect matchings in a planar graph, a matchgate-based solution might not necesssarily exist because the range of encodable local constraints is too limited. Valiant's insight in [21, 22] was to extend this range by looking instead for a *many-to-many* reduction. In his new framework, multiple strands of computation get combined in a "holographic" mixture with exponential, custom-built cancellations specified by a choice of basis vectors to produce the final answer. With this change-of-basis approach, Valiant [21, 22, 23, 24] found polynomial-time solutions to a number of counting problems,

*Harvard University Department of Mathematics & Harvard John A. Paulson School of Engineering and Applied Sciences

minor variants of which are known to be intractable, and noted that the only criterion for their existence was the solvability of certain finite systems of polynomial equations. As a notable example, whereas counting the number of satisfying assignments to planar, read-twice, monotone 3-CNFs is #P-complete and even the same problem modulo 2 is NP-hard under randomized reductions, the same problem modulo 7, known as $\#_7$Pl-Rtw-Mon-3CNF, has a polynomial-time holographic solution [21].

Given the dearth of strong unconditional circuit lower bounds to support the prevailing belief that $\mathsf{P} \neq \mathsf{NP}$, we might question whether our current inability to devise polynomial-time solutions to NP-complete problems is sufficient to justify such a belief. After all, it would also have justified the belief that $\#_7$Pl-Rtw-Mon-3CNF is intractable. This suggests that to arrive upon the desired separation of P and NP, we need a better understanding of the possibilities of polynomial-time computation. For this reason, determining the ultimate capabilities of holographic algorithms appears to be a crucial step.

To this end, one major direction has been to study the kinds of local constraints that matchgates can encode. Cai and his collaborators initiated a systematic study of the structural theory of holographic algorithms over the Boolean domain [1, 2, 3, 4, 5, 6, 7, 8, 9, 10], compiling what amounts to a catalogue of such constructions that turns the process of finding basic holographic reductions into something essentially algorithmic. For example, as a corollary to their results in [7], they noted that the modulus 7 appears in $\#_7$Pl-Rtw-Mon-3CNF because it is a Mersenne prime and that more generally, there is a polynomial-time holographic algorithm for $\#_{2^k-1}$Pl-Rtw-Mon-kCNF.

1.2 Basis Collapse Theorems

Another direction of study concerns better understanding the power afforded by the change of basis, specifically the relationship between two important parameters of holographic algorithms: the dimension and number of basis vectors. Because holographic algorithms reduce counting problems to counting perfect matchings, the dimension of the k basis vectors specifying the abovementioned "holographic mixture" must be of the form 2^ℓ. We say that ℓ is the *basis size* and k is the domain size. The domain size should be interpreted as the range over which variables in the counting problem can take values, so for instance, problems related to counting certain k-colorings in a graph are problems over domain size k. The basis size should then be interpreted roughly as the number of bits needed to encode each of these k colors.

The first holographic algorithms studied were on domain size 2, dealing with counting problems involving matchings, 2-colorings, graph bipartitions, and Boolean satisfying assignments, and used bases of size 1. The original holographic algorithm for $\#_7$Pl-Rtw-Mon-3CNF was the first to use a basis of size 2, and at the time it may have been tempting to believe that increasing basis size adds power to the holographic approach. In [8], Cai and Lu showed to the contrary that on domain size 2, any nontrivial holographic algorithm with basis of size $\ell \geq 2$ can be simulated by one with basis of size 1. In other words, for all ℓ, the class of domain size 2 problems solvable with basis size ℓ *collapses* to the class solvable with basis size 1.

In [5], Cai and Fu further showed that holographic algorithms on domain sizes 3 and 4 using at least one full-rank signature collapse to basis sizes 1 and 2 respectively. They conjectured that for domain size k, we get a collapse to basis size $\lfloor \log_2 k \rfloor$ and suggested a heuristic explanation that for domain size $k = 2^K$, we only need $\log_2 k$ bits to encode each of the k colors.

1.3 Our Results and Techniques

We prove Cai and Fu's conjecture in the affirmative for all domain sizes.

Theorem 1.1. *Any holographic algorithm on domain size k using at least one matchgate with signature of rank k can be simulated by a holographic algorithm with basis size $\lfloor \log_2 k \rfloor$.*

As Cai and Fu noted in [5], their "information-theoretic" explanation for the collapse theorems on domain sizes up to 4 is insufficient to explain why holographic algorithms on domain size 3 collapse to basis size 1 and not just 2.

To prove their collapse theorem on domain size 3, Cai and Fu actually showed that the bases of holographic algorithms on domain size 3 which use at least one full-rank signature must be of rank 2 rather than 3. They then invoked the main result of Fu and Yang [12] which reduces holographic algorithms on domain size k with bases of rank 2 to ones on domain size 2.

Our key observation is that this phenomenon occurs at a much larger scale. As a bit of informal background, the *standard signature* of a matchgate G is a vector encoding the perfect matching properties of G. By indexing appropriately, we can regard the standard signature as a matrix Γ. The entries of this matrix are known [6] to satisfy a collection of quadratic polynomial identities called the Matchgate Identities (MGIs), and by using these identities together with some multilinear algebra, we prove the following:

Theorem 1.2 (Rank Rigidity Theorem)**.** *The rank of the standard signature matrix Γ is always a power of two.*

We can then conclude that the basis of a nontrivial holographic algorithm on domain size k must be of rank 2^ℓ, where ℓ is the largest integer for which $2^\ell \leq k$. With this step, together with a generalization of Fu and Yang's result to bases of rank k, we show it is enough to prove a collapse theorem for holographic algorithms on domain sizes that are powers of two. Cai and Fu [5] achieved such a collapse theorem for domain size 4 by proving that 1) any standard signature of rank 4 contains a full-rank 4×4 submatrix whose entries have indices are "close" in Hamming distance, 2) full-rank $4 \times 2^{2n-2}$ standard signatures have right inverses that are also standard signatures.

For 1), the proof in [5] used algebraic techniques involving the matchgate identities, but these methods seem to work only for domain size 4. We instead show that the required generalization of 1) to arbitrary domain sizes almost trivially follows from the rank rigidity theorem and the MGIs. Roughly, we prove the following:

Theorem 1.3 (Cluster Existence - informal)**.** *If Γ is a $2^\ell \times 2^{(n-1)\ell}$ matrix of rank at least k realizable as the standard signature of some matchgate, then there exists a $2^{\lceil \log_2 k \rceil} \times 2^{\lceil \log_2 k \rceil}$ submatrix of full rank in Γ whose column (resp. row) indices differ in at most $\lceil \log_2 k \rceil$ bits.*

For 2), the proof in [5] nonconstructively verifies that the set of all invertible 4×4 matrices satisfying the matchgate identities up to sign forms a group under multiplication. 4×4 matrices are easy to handle because there is only one nontrivial MGI in this case. Rather than generalizing this approach, we note that Li and Xia [15] proved a very similar but more general result under a different framework of matchgate computation known as character theory, showing that the set of all invertible $2^K \times 2^K$ matrices realizable as matchgate characters forms a group under multiplication for all K. It turns out their technique carries over with minor modifications into the framework of signature theory that we consider, and we use it to show the following:

Theorem 1.4 (Group Property - informal)**.** *If G is a generator matchgate with $2^K \times 2^{(n-1)K}$ standard signature $\underline{G}$ of rank 2^K, then there exists a recognizer matchgate with $2^{K(n-1)} \times 2^K$ standard signature $\underline{R}$ such that $\underline{G} \cdot \underline{R} = I_{2^K}$, where I_{2^K} denotes the $2^K \times 2^K$ identity matrix.*

Our general collapse theorem then follows from our generalizations of 1) and 2). This result gives a way forward for the development of the structural theory of holographic algorithms on higher

domain sizes in the same vein as Cai et al.'s work on domain size 2. In [24], Valiant gave examples of holographic algorithms on domain size 3, but holographic algorithms on higher domain sizes have yet to be explored. Our result shows that for domain size k, we can focus on understanding changes of basis in $\mathcal{M}_{2^{\lfloor \log_2 k \rfloor} \times k}(\mathbb{C})$ rather than over an infinite set of dimensions, just as the collapse theorem of Cai and Lu [8] showed that on the Boolean domain, they could focus on understanding changes of basis in $\mathrm{GL}_2(\mathbb{C})$.

1.4 Organization

In Section 2, we introduce the necessary notation and background, including a brief overview of holographic algorithms. In Section 3, we state the matchgate identities and some of their key consequences. In Section 4, we prove Theorems 1.2 and 1.3. In Section 5, we then prove Theorem 1.4. In Section 6, we generalize the main result of Fu and Yang [12] to bases of rank k and reduce proving a collapse theorem on all domain sizes to proving one on domain sizes equal to powers of two. Finally, in Section 7, we prove the desired collapse theorem, Theorem 1.1, on domain sizes equal to powers of two by invoking the results from Section 4.

2 Preliminaries

2.1 Background

Denote the Hamming weight of string α by $\mathrm{wt}(\alpha)$, and define the parity of α to be the parity of $\mathrm{wt}(\alpha)$. Given $1 \leq i \leq m$, define $e_i \in \{0,1\}^m$ to be the bitstring with a single nonzero bit in position i. The parameter m is implicit, and when this notation is used, m will be clear from the context. Denote by 1^m the length-m bitstring consisting solely of 1's.

We review some basic definitions and results about holographic algorithms. For a comprehensive introduction to this subject, see [22].

Definition 2.1. A *matchgate* $\Gamma = (G, X, Y)$ is defined by a planar embedding of a planar graph $G = (V, E, W)$, *input nodes* $X \subseteq V$, and *output nodes* $Y \subseteq V$, where $X \cap Y = \emptyset$. We refer to $X \cup Y$ as the *external nodes* of Γ.

We say that Γ has *arity* $|X| + |Y|$. In the planar embedding of G, the input and output nodes are arranged such that if one travels counterclockwise around the outer face of G, one encounters first the input nodes labeled 1,2,...,$|X|$ and then the output nodes $|Y|$,...,2,1.

If Γ has exclusively output (resp. input) nodes, we say that Γ is a generator (resp. recognizer). Otherwise, we say that Γ is an $|X|$-*input,* $|Y|$-*output transducer.*

Definition 2.2. A *basis* of size ℓ on domain size k is a $2^\ell \times k$ matrix $M = (a_i^\alpha)$, where rows and columns are indexed by $\alpha \in \{0,1\}^\ell$ and $i \in [k]$ respectively.

Definition 2.3. The *standard signature* of a matchgate Γ of arity $n\ell$ is a vector of dimension $2^{n\ell}$ which will be denoted by $\underline{\Gamma}$, where for $\alpha_i \in \{0,1\}^\ell$, $\underline{\Gamma}^{\alpha_1 \cdots \alpha_n}$ denotes the entry of $\underline{\Gamma}$ indexed by $\alpha_1 \circ \cdots \circ \alpha_n$. If Z is the subset of the external nodes of Γ for which $\alpha_1 \circ \cdots \circ \alpha_n$ is the indicator string, then

$$\underline{\Gamma}^{\alpha_1 \cdots \alpha_n} = \mathrm{PerfMatch}(\Gamma \backslash Z).$$

Here, if $A = (A_{ij})$ is the adjacency matrix of Γ, PerfMatch is the polynomial in the entries of A defined by

$$\mathrm{PerfMatch}(A) = \sum_M \prod_{(i,j) \in M} A_{ij},$$

with the sum taken over the set M of all perfect matchings of Γ.

The following lemma follows from the definition of standard signatures.

Lemma 2.4. *Suppose $\underline{R}$ is the standard signature of a recognizer of arity $n\ell$ and T the standard signature of a transducer with s inputs and ℓ outputs. Then $\underline{R}' = \underline{R}T^{\otimes n}$ is the standard signature of a recognizer matchgate of arity ns.*

Definition 2.5. A column (resp. row) vector of dimension k^n is said to be a *generator (resp. recognizer) signature realizable over a basis M* if there exists a generator (resp. recognizer) matchgate Γ satisfying $M^{\otimes n}G = \underline{G}$ (resp. $\underline{R}M^{\otimes n} = R$). We say that a collection of recognizer and generator signatures $R_1, ..., R_a, G_1, ..., G_b$ is *simultaneously realizable* if they are realizable over a common basis M.

In particular, if M is square, the signature of a matchgate with respect to the standard basis is the standard signature. Also note that in terms of coordinates, we have that

$$\underline{G}^{\alpha_1 \cdots \alpha_n} = \sum_{j_1, ..., j_n \in [k]} G^{j_1 \cdots j_n} a_{j_1}^{\alpha_1} \cdots a_{j_n}^{\alpha_n}$$

and

$$R_{j_1 \cdots j_n} = \sum_{\alpha_1, ..., \alpha_n \in \{0,1\}^\ell} \underline{R}_{\alpha_1 \cdots \alpha_n} a_{j_1}^{\alpha_1} \cdots a_{j_n}^{\alpha_n}.$$

Definition 2.6. A *matchgrid* $\Omega = (G, R, W)$ is a weighted planar graph consisting of a set of g generators $G = \{G_1, ..., G_g\}$, a set of r recognizers $R = \{R_1, ..., R_r\}$, and a set of w wires $W = \{W_1, ..., W_w\}$, each of which has weight 1 and connects the output node of a generator to the input node of a recognizer so that every input and output node among the matchgates in $G \cup R$ lies on exactly one wire.

Definition 2.7. Suppose $\Omega = (G, R, W)$ is a matchgrid with g generators, r recognizers, and w wires, and let M be a basis for Ω. Define the *Holant* to be the following quantity:

$$\text{Holant}(\Omega) = \sum_{z \in [k]^w} \left(\prod_{i=1}^{g} G_i^{y_i} \prod_{j=1}^{r} R_j^{x_j} \right).$$

Here, $z = y_1 \circ \cdots \circ y_g = x_1 \circ \cdots \circ x_r$ such that $y_i \in [k]^{|Y_i|}$ and $x_j \in [k]^{|X_j|}$ for Y_i the output nodes of G_i and X_j the input nodes of R_j, and G_i and R_j denote the signatures of their respective matchgates under basis M.

Valiant's Holant theorem states the following.

Theorem 2.8 (Theorem 4.1, [22])**.** *If Ω is a matchgrid over a basis M, then* $\text{Holant}(\Omega) = \text{PerfMatch}(\Omega)$.

As the Fisher-Kasteleyn-Temperley algorithm [13, 17] can compute the number of perfect matchings of a planar graph in polynomial time, $\text{Holant}(\Omega)$ can be computed in polynomial time as long as Ω is planar.

2.2 Matrix Form of Signatures

It will be convenient to regard signatures not as vectors but as matrices.

Definition 2.9. For generator signature G, the t-th matrix form $G(t)$ $(1 \leq t \leq n)$ is a $k \times k^{n-1}$ matrix where the rows are indexed by $1 \leq j_t \leq k$ and the columns are indexed by $j_1 \cdots j_{t-1} j_{t+1} \cdots j_n$ in lexicographic order.

Definition 2.10. For recognizer signature R, the t-th matrix form $R(t)$ $(1 \leq t \leq n)$ is a $k^{n-1} \times k$ matrix where the rows are indexed by $j_1 \cdots j_{t-1} j_{t+1} \cdots j_n$ in lexicographic order and the columns are indexed by $1 \leq j_t \leq k$.

We would also like to regard standard signatures as matrices; if basis M is square, the following definitions are special cases of the above.

Definition 2.11. For standard signature $\underline{G}$, the t-th matrix form $\underline{G}(t)$ $(1 \leq t \leq n)$ is a $2^\ell \times 2^{(n-1)\ell}$ matrix where the rows are indexed by α_t and the columns are indexed by $\alpha_1 \cdots \alpha_{t-1} \alpha_{t+1} \cdots \alpha_n$.

Definition 2.12. For standard signature $\underline{R}$, the t-th matrix form $\underline{R}(t)$ $(1 \leq t \leq n)$ is a $2^{(n-1)\ell} \times 2^\ell$ matrix where the rows are indexed by $\alpha_1 \cdots \alpha_{t-1} \alpha_{t+1} \cdots \alpha_n$ and the columns are indexed by α_t.

One can check that $\underline{G}(t)$ and $G(t)$, and $\underline{R}(t)$ and $R(t)$, are related as follows.

Lemma 2.13. *If* $\underline{G} = M^{\otimes n} G$, *then*

$$\underline{G}(t) = M G(t) (M^T)^{\otimes (n-1)}.$$

Lemma 2.14. *If* $R = \underline{R} M^{\otimes n}$, *then*

$$R(t) = (M^T)^{\otimes (n-1)} \underline{R}(t) M.$$

We will denote by $\underline{G}(t)^\sigma$ the row vector indexed by σ, $\underline{G}(t)_\zeta$ the column vector indexed by ζ, and $\underline{G}(t)^\sigma_\zeta$ the entry of $\underline{G}$ in row σ and column ζ. We use analogous notation for matrices $\underline{R}$, G, and R. In general, if Γ is any matrix, we will sometimes refer to the entry Γ^σ_ζ as the "entry (indexed by) (σ, ζ)."

In general, if Γ is a matrix with rows indexed by $\{0,1\}^a$ and columns indexed by $\{0,1\}^b$, and $S \subset \{0,1\}^a$ (resp. $S \subset \{0,1\}^b$), we will let Γ^S (resp. Γ_S) denote the submatrix of Γ consisting of rows (resp. columns) indexed by S. Where Γ is clear from context, we will denote the row span of Γ^S (resp. column span of Γ_S) by $\mathrm{span}(S)$.

Lastly, a column/row is called *odd* (resp. *even*) if its index is odd (resp. even).

2.3 Degenerate and Full Rank Signatures

Definition 2.15. A signature G (generator or recognizer) is degenerate iff there exist vectors γ_i $(1 \leq i \leq n)$ of dimension k for which $G = \gamma_1 \otimes \cdots \otimes \gamma_n$.

Lemma 2.16 (Lemma 3.1, [5])**.** *A signature G is degenerate iff* $\mathrm{rank}(G(t)) \leq 1$ *for* $1 \leq t \leq n$.

Definition 2.17. A signature G is of full rank iff there exists some $1 \leq t \leq n$ for which $\mathrm{rank}(G(t)) = k$.

By Lemma 2.13, if signature G is of full rank, then for the corresponding standard signature $\underline{G}$, we have that $\text{rank}(\underline{G}(t)) = k$ for some t. Over domain size 2, by Lemma 2.16, all signatures not of full rank are degenerate, and holographic algorithms exclusively using such signatures are trivial because degenerate generators can by definition be decoupled into arity-1 generators. Over domain size $k \geq 3$ however, it is unknown to what extent holographic algorithms exclusively using signatures not of full rank trivialize. In [5], the collapse theorems over domain sizes 3 and 4 were proved under the assumption that at least one signature is of full rank, so we too make that assumption.

2.4 Clusters

One of the key results in our proof of the general collapse theorem is the existence within any matrix-form standard signature of a full-rank square submatrix whose entries have indices satisfying certain properties. In this section we make precise those properties.

Definition 2.18. A set of 2^m distinct bitstrings $Z = \{x^1, ..., x^{2^m}\} \subset \{0,1\}^n$ is an (m,n)-*cluster* if there exists $s \in \{0,1\}^n$ and positions $p_1, ..., p_m \in [n]$ such that for each $i \in [2^m]$, $x^i = s \oplus \left(\bigoplus_{j \in J} e_{p_j}\right)$ for some $J \subset \{p_1, ..., p_m\}$. We write Z as $s + \{e_{p_1}, ..., e_{p_m}\}$. Note that for a fixed cluster, s is only unique up to the bits outside of positions $p_1, ..., p_m$. If a cluster Z' is a subset of another cluster Z, we say that Z' is a *subcluster* of Z.

Definition 2.19. In Γ, a $2^m \times 2^m$ submatrix Γ' is a m-*cluster submatrix* if there exist (m,ℓ)- and $(m,(n-1)\ell)$-clusters Σ and Z such that $\Gamma' = (\Gamma^{\sigma}_{\zeta})^{\sigma \in \Sigma}_{\zeta \in Z}$ (here we omit the parameters n and ℓ in the notation as they will be clear from context).

3 Matchgate Identities and Consequences

3.1 Parity Condition and Matchgate Identities

The most obvious property that standard signatures satisfy is the *parity condition*: because a graph with an odd number of vertices has no perfect matchings, the indices of the nonzero entries in $\underline{G}$ have the same parity. It trivially follows that in $\underline{G}(t)$, columns $\underline{G}(t)_{\zeta}$ and $\underline{G}(t)_{\eta}$, if both nonzero, are linearly independent if ζ and η are of opposite parities. The same holds for the rows of $\underline{G}(t)$.

In [6] it is shown algebraically that the parity condition in fact follows from the so-called *matchgate identities* which we present below, quadratic relations which together form a necessary and sufficient condition for a vector to be the standard signature of some matchgate.

As in Cai and Fu's proof of the collapse theorem for domain size 4, we will make heavy use of the matchgate identities stated below. Wherever we invoke them, they will be for generator matchgates, so we focus on this case.

Theorem 3.1 (Theorem 2.1, [6])**.** *A $2^{\ell} \times 2^{(n-1)\ell}$ matrix Γ is the t-th matrix form of the standard signature of some generator matchgate iff for all $\zeta, \eta \in \{0,1\}^{(n-1)\ell}$ and $\sigma, \tau \in \{0,1\}^{\ell}$, the following* matchgate identity *holds. Let $\zeta \oplus \eta = e_{q_1} \oplus \cdots \oplus e_{q_{d'}}$ and $\sigma \oplus \tau = e_{p_1} \oplus \cdots \oplus e_{p_d}$, where $q_1 < \cdots < q_{d'}$ and $p_1 < \cdots < p_d$. Then*

$$\sum_{i=1}^{d} (-1)^{i+1} \Gamma^{(\sigma \oplus e_{p_1} \oplus e_{p_i})}_{\zeta} \Gamma^{(\tau \oplus e_{p_1} \oplus e_{p_i})}_{\eta} = \sum_{j=1}^{d'} \pm \Gamma^{(\sigma \oplus e_{p_1})}_{(\zeta \oplus e_{q_j})} \Gamma^{(\tau \oplus e_{p_1})}_{(\eta \oplus e_{q_j})}. \tag{1}$$

Here the $\pm$ signs depend on both j and, if q_j is after the t-th block, the parity of d. If d is even,

$$\sum_{i=1}^{d}(-1)^{i+1}\Gamma_{\zeta}^{(\sigma\oplus e_{p_1}\oplus e_{p_i})}\Gamma_{\eta}^{(\tau\oplus e_{p_1}\oplus e_{p_i})} = \epsilon_{\zeta,\eta}\sum_{j=1}^{d'}(-1)^{j+1}\Gamma_{(\zeta\oplus e_{q_j})}^{(\sigma\oplus e_{p_1})}\Gamma_{(\eta\oplus e_{q_j})}^{(\tau\oplus e_{p_1})}, \tag{2}$$

where $\epsilon_{\zeta,\eta}\in\{\pm 1\}$ is positive (resp. negative) if the number of q_j preceding the t-th block is odd (resp. even).

Remark 3.2. *If $d+d'$ is odd,* (1) *is trivial by the parity condition.*

We will be making extensive use of the matchgate identities in this paper, but we will typically not care about the $\epsilon_{\zeta,\eta}$ sign on the right-hand side of (2). For this reason, it will be convenient to make the following definition.

Definition 3.3. A $2^{\ell}\times 2^{m}$ matrix M is a *pseudo-signature* if for all σ,τ for which $\mathrm{wt}(\sigma\oplus\tau)$ is even, its entries satisfy the corresponding identity (2) up to a factor of ± 1 on the right-hand side.

Standard signatures and cluster submatrices are all examples of pseudo-signatures.

Observation 1. *If M is a pseudo-signature, then its transpose M^T is a pseudo-signature.*

3.2 Matchgate Identities and Determinants

We now derive from the matchgate identities some basic linear algebraic properties of the columns of pseudo-signatures. By Observation 1, these also apply for the rows.

Firstly, we have the following immediate consequence of Theorem 3.1. We specifically consider the case where $\mathrm{wt}(\zeta\oplus\eta)$ is even and, by Remark 3.2, $\mathrm{wt}(\sigma\oplus\tau)$ is even. So write $\zeta\oplus\eta = e_{q_1}\oplus\cdots\oplus e_{q_{2d'}}$ and $\sigma\oplus\tau = e_{p_1}\oplus\cdots\oplus e_{p_{2d}}$.

Reverse the roles of ζ and η in (2). Subtract the resulting equation from (2) to find

$$\sum_{i=1}^{2d}(-1)^{i+1}\left(\Gamma_{\eta}^{(\sigma\oplus e_{p_1}\oplus e_{p_i})}\Gamma_{\zeta}^{(\tau\oplus e_{p_1}\oplus e_{p_i})} - \Gamma_{\zeta}^{(\sigma\oplus e_{p_1}\oplus e_{p_i})}\Gamma_{\eta}^{(\tau\oplus e_{p_1}\oplus e_{p_i})}\right) =$$

$$\epsilon_{\zeta,\eta}\sum_{j=1}^{2d'}(-1)^{j+1}\left(\Gamma_{(\zeta\oplus e_{q_j})}^{(\sigma\oplus e_{p_1})}\Gamma_{(\eta\oplus e_{q_j})}^{(\tau\oplus e_{p_1})} - \Gamma_{(\eta\oplus e_{q_j})}^{(\sigma\oplus e_{p_1})}\Gamma_{(\zeta\oplus e_{q_j})}^{(\tau\oplus e_{p_1})}\right),$$

or equivalently,

$$\sum_{i=1}^{2d}(-1)^{i+1}\begin{vmatrix}\Gamma_{\zeta}^{(\sigma\oplus e_{p_1}\oplus e_{p_i})} & \Gamma_{\eta}^{(\sigma\oplus e_{p_1}\oplus e_{p_i})}\\ \Gamma_{\zeta}^{(\tau\oplus e_{p_1}\oplus e_{p_i})} & \Gamma_{\eta}^{(\tau\oplus e_{p_1}\oplus e_{p_i})}\end{vmatrix} = \epsilon_{\zeta,\eta}\sum_{j=1}^{2d'}(-1)^{j+1}\begin{vmatrix}\Gamma_{(\zeta\oplus e_{q_j})}^{(\sigma\oplus e_{p_1})} & \Gamma_{(\eta\oplus e_{q_j})}^{(\sigma\oplus e_{p_1})}\\ \Gamma_{(\zeta\oplus e_{q_j})}^{(\tau\oplus e_{p_1})} & \Gamma_{(\eta\oplus e_{q_j})}^{(\tau\oplus e_{p_1})}\end{vmatrix} \tag{3}$$

Example 3.4. *Suppose $d=d'=1$. Then* (3) *becomes*

$$2\begin{vmatrix}\Gamma_{\zeta}^{\sigma} & \Gamma_{\eta}^{\sigma}\\ \Gamma_{\zeta}^{\tau} & \Gamma_{\eta}^{\tau}\end{vmatrix} = 2\epsilon_{\zeta,\eta}\begin{vmatrix}\Gamma_{\zeta\oplus e_{q_1}}^{\sigma\oplus e_{p_1}} & \Gamma_{\zeta\oplus e_{q_2}}^{\sigma\oplus e_{p_1}}\\ \Gamma_{\zeta\oplus e_{q_1}}^{\sigma\oplus e_{p_2}} & \Gamma_{\zeta\oplus e_{q_2}}^{\sigma\oplus e_{p_2}}\end{vmatrix},$$

so in particular, the matrix on the left is singular iff the latter is.

More generally, only suppose that $d = 1$. Then (3) *becomes*

$$2\begin{vmatrix}\Gamma^{\sigma}_{\zeta} & \Gamma^{\sigma}_{\eta}\\ \Gamma^{\tau}_{\zeta} & \Gamma^{\tau}_{\eta}\end{vmatrix} = \epsilon_{\zeta,\eta}\sum_{j=1}^{2d'}(-1)^{j+1}\begin{vmatrix}\Gamma^{(\sigma\oplus e_{p_1})}_{(\zeta\oplus e_{q_j})} & \Gamma^{(\sigma\oplus e_{p_1})}_{(\eta\oplus e_{q_j})}\\ \Gamma^{(\sigma\oplus e_{p_2})}_{(\zeta\oplus e_{q_j})} & \Gamma^{(\sigma\oplus e_{p_2})}_{(\eta\oplus e_{q_j})}\end{vmatrix} \tag{4}$$

so in particular, the matrix on the left-hand side is singular if all $2d'$ matrices on the right-hand side are singular.

In other words, if $\ell = 2$ so that Γ only has four rows, columns ζ and η as defined above are linearly dependent if all pairs of neighboring columns are linearly dependent. We shall see in the next section (Corollary 3.6) that this is true even when Γ has an arbitrary number of rows.

3.3 Wedge Products of Columns

Motivated by Example 3.4, we'd like to study the set of all 2×2 determinants $\begin{vmatrix}\Gamma^{\sigma}_{\zeta} & \Gamma^{\sigma}_{\eta}\\ \Gamma^{\tau}_{\zeta} & \Gamma^{\tau}_{\eta}\end{vmatrix}$ given two column vectors Γ_ζ and Γ_η of the same parity. These are merely the coefficients of the wedge product $\Gamma_\zeta \wedge \Gamma_\eta$ under the standard basis $\{v_\sigma \wedge v_\tau\}_{\sigma,\tau\in\{0,1\}^{2^\ell},\ \sigma<\tau}$ (where the relation $\sigma < \tau$ denotes lexicographic ordering) of $\bigwedge^2 \mathbb{C}^{2^\ell}$, the second exterior power of $\mathbb{C}^{2^\ell}$. The matchgate identities imply the following consequence about the relationships among the wedge products $\Gamma_\zeta \wedge \Gamma_\eta$ as ζ and η vary.

Lemma 3.5. *If $\zeta_1, \eta_1, \zeta_2, \eta_2, ..., \zeta_m, \eta_m$ are even indices for which*

$$\sum_{\nu=1}^{m} a_\nu \left(\Gamma_{\zeta_\nu} \wedge \Gamma_{\eta_\nu}\right) = 0 \tag{5}$$

for some $a_1, ..., a_m \in \mathbb{C}$, then

$$\sum_{\nu=1}^{m} \epsilon_{\zeta_\nu,\eta_\nu} a_\nu \left(\sum_{i=1}^{2d_\nu}(-1)^{j+1}\Gamma_{\zeta_\nu\oplus e_{p_i^\nu}} \wedge \Gamma_{\eta_\nu\oplus e_{p_i^\nu}}\right) = 0, \tag{6}$$

where for each ν, $\mathrm{wt}(\zeta_\nu \oplus \eta_\nu) = 2d_\nu$ and $\zeta_\nu \oplus \eta_\nu = e_{p_1^\nu} \oplus \cdots \oplus e_{p_{2d_\nu}^\nu}$.

Proof. For convenience, we will denote $\epsilon_{\zeta_\nu,\eta_\nu}$ by ϵ_ν. First, we rewrite (6) in terms of coordinates as

$$\sum_{\nu=1}^{m} \epsilon_\nu a_\nu \left(\sum_{i=1}^{2d_\nu}(-1)^{i+1}\sum_{\sigma<\tau}\begin{vmatrix}\Gamma^{\sigma}_{\zeta_\nu\oplus e_{p_i^\nu}} & \Gamma^{\sigma}_{\eta_\nu\oplus e_{p_i^\nu}}\\ \Gamma^{\tau}_{\zeta_\nu\oplus e_{p_i^\nu}} & \Gamma^{\tau}_{\eta_\nu\oplus e_{p_i^\nu}}\end{vmatrix}(v_\sigma \wedge v_\tau)\right) = 0, \tag{7}$$

where $\sigma < \tau$ denotes lexicographical ordering. Note that the determinants that appear in the left-hand side of (7) are zero when σ and τ are of opposite parity. Moreover, depending on the parity of the signature Γ, either all such determinants are also zero for σ and τ even, or they are all zero for σ and τ odd.

Rearranging the order of summations in (7), the desired identity becomes

$$\sum_{\sigma<\tau}(v_\sigma \wedge v_\tau)\cdot\left(\sum_{\nu=1}^{m}\epsilon_\nu a_\nu \sum_{i=1}^{2d_\nu}(-1)^{i+1}\begin{vmatrix}\Gamma^{\sigma}_{\zeta_\nu\oplus e_{p_i^\nu}} & \Gamma^{\sigma}_{\eta_\nu\oplus e_{p_i^\nu}}\\ \Gamma^{\tau}_{\zeta_\nu\oplus e_{p_i^\nu}} & \Gamma^{\tau}_{\eta_\nu\oplus e_{p_i^\nu}}\end{vmatrix}\right) = 0. \tag{8}$$

For a fixed pair $\sigma < \tau$, let $\mathrm{wt}(\sigma \oplus \tau) = 2d'$ and $\sigma \oplus \tau = e_{q_1} \oplus \cdots \oplus e_{q_{2d'}}$. If we apply (3) and rearrange the order of summation once more, the coefficient of $v_\sigma \wedge v_\tau$ above becomes

$$\sum_{\nu=1}^{m} a_\nu \cdot \sum_{j=1}^{2d'} (-1)^{j+1} \begin{vmatrix} \Gamma_{\zeta_\nu}^{\sigma \oplus e_{q_j}} & \Gamma_{\eta_\nu}^{\sigma \oplus e_{q_j}} \\ \Gamma_{\zeta_\nu}^{\tau \oplus e_{q_j}} & \Gamma_{\eta_\nu}^{\tau \oplus e_{q_j}} \end{vmatrix} = \sum_{j=1}^{2d'} (-1)^{j+1} \sum_{\nu=1}^{m} a_\nu \begin{vmatrix} \Gamma_{\zeta_\nu}^{\sigma \oplus e_{q_j}} & \Gamma_{\eta_\nu}^{\sigma \oplus e_{q_j}} \\ \Gamma_{\zeta_\nu}^{\tau \oplus e_{q_j}} & \Gamma_{\eta_\nu}^{\tau \oplus e_{q_j}} \end{vmatrix}.$$

But note that if we expand (5) in terms of coordinates, the term

$$a_\nu \begin{vmatrix} \Gamma_{\zeta_\nu}^{\sigma \oplus e_{q_j}} & \Gamma_{\eta_\nu}^{\sigma \oplus e_{q_j}} \\ \Gamma_{\zeta_\nu}^{\tau \oplus e_{q_j}} & \Gamma_{\eta_\nu}^{\tau \oplus e_{q_j}} \end{vmatrix}$$

is precisely the coefficient of $v_{\sigma \oplus e_{q_j}} \wedge v_{\tau \oplus e_{q_j}}$ in the expansion of (5) in terms of coordinates and hence zero by assumption, so (8) holds as desired. □

Corollary 3.6. *Let $\zeta, \eta \in \{0,1\}^{(n-1)\ell}$ be such that $\zeta \oplus \eta = \bigoplus_{j=1}^{2d} e_{p_i}$. If column $\Gamma_{(\zeta \oplus e_{p_i})}$ is linearly dependent with column $\Gamma_{(\eta \oplus e_{p_i})}$ for $1 \le i \le 2d$, then column Γ_ζ is linearly dependent with column Γ_η.*

Corollary 3.7. *Let $\zeta, \eta \in \{0,1\}^{(n-1)\ell}$ be such that $\zeta \oplus \eta = \bigoplus_{i=1}^{2d} e_{p_i}$. If there exists $i \in [2d']$ such that column $\Gamma_{\zeta \oplus e_{p_i}}$ is linearly dependent with column $\Gamma_{\eta \oplus e_{p_i}}$ for $i = 1, ..., \hat{j}, ..., 2d'$, where $\hat{j}$ denotes omission of index j, and if Γ_ζ is also linearly dependent with column Γ_η, then $\Gamma_{\zeta \oplus e_{p_j}}$ and $\Gamma_{\eta \oplus e_{p_j}}$ are linearly dependent.*

Lemma 3.5 says that any linear relation among wedges of even columns yields a linear relation among wedges of odd columns, and vice versa.

Lastly, we need the following elementary result in multilinear algebra.

Lemma 3.8. *If $v_1, ..., v_n$ are linearly independent in vector space V, then the set of all $v_i \wedge v_j$ for $i < j$ are linearly independent in $\bigwedge^2 V$.*

Combining this with Lemma 3.5 yields the following key ingredient to the analysis in Section 4.

Lemma 3.9. *Suppose $\zeta_0, \eta \in \{0,1\}^K$ such that $\zeta_0 \neq \eta$, and the indices in $T = \{\zeta_1, ..., \zeta_m\} \subset \{0,1\}^K$ are distinct and have the same parity. Suppose further that $\zeta_0 \neq \zeta_1, ..., \zeta_m$. Let $\zeta_i \oplus \eta = e_{p_1^i} \oplus \cdots \oplus e_{p_{d_i}^i}$ for $0 \le i \le m$, where $d_i := \mathrm{wt}(\zeta_i \oplus \eta)$. Define*

$$S := \bigcup_{0 \le i \le m, 1 \le j \le d_i} \{\eta \oplus e_{p_j^i}, \zeta_i \oplus e_{p_j^i}\} \subset \{0,1\}^K$$

not in the sense of multisets, that is, we throw out duplicates so that the strings in S are all distinct.

Suppose the columns indexed by S are linearly independent. Then

$$\Gamma_{\zeta_0} \notin \mathit{span}(\Gamma_{\zeta_1}, ..., \Gamma_{\zeta_m}) \tag{9}$$

If $\mathrm{wt}(\eta \oplus \zeta_0) \ge 4$ and $j^* \in [d_0]$, then (9) holds even if only the columns indexed by $S' := S \backslash \{\zeta_0 \oplus e_{p_{j^*}^0}\}$ are linearly independent.

Proof. We first prove the claim without the assumption that $\mathrm{wt}(\eta \oplus \zeta_0) \geq 4$. Suppose to the contrary that $\Gamma_{\zeta_0} = \sum_{i=1}^{m} a_i \Gamma_{\zeta_i}$ so that $\Gamma_{\zeta_0} \wedge \Gamma_\eta - \sum_{i=1}^{m} a_i \Gamma_{\zeta_i} \wedge \Gamma_\eta = 0$. By Lemma 3.5, this linear relation implies the following linear relation among wedges of columns of the other parity:

$$\epsilon_{\zeta_0,\eta} \left(\sum_{j=1}^{d_0} (-1)^{j+1} \Gamma_{\zeta_0 \oplus e_{p_j^0}} \wedge \Gamma_{\eta \oplus e_{p_j^0}} \right) - \sum_{i=1}^{m} \epsilon_{\zeta_i,\eta} a_i \left(\sum_{j=1}^{d_i} (-1)^{j+1} \Gamma_{\zeta_i \oplus e_{p_j^i}} \wedge \Gamma_{\eta \oplus e_{p_j^i}} \right) = 0 \tag{10}$$

We claim this is a nontrivial linear relation contradicting the linear independence of the columns indexed by S. For each of the $m+1$ sums indexed by $1 \leq j \leq d_i$ appearing in (10), if $d_i = 2$, rewrite $\sum_{j=1}^{d_i} (-1)^{j+1} \Gamma_{\zeta_i \oplus e_{p_j^i}} \wedge \Gamma_{\eta \oplus e_{p_j^i}}$ as $2\Gamma_{\zeta_i \oplus e_{p_1^i}} \wedge \Gamma_{\eta \oplus e_{p_1^i}}$.

After this consolidation, note that the wedge products in (10) are now all distinct. Certainly for any $j, j' \in [d_i]$ where $d_i \geq 4$, $\Gamma_{\zeta_i \oplus e_{p_j^i}} \wedge \Gamma_{\eta \oplus e_{p_j^i}}$ and $\Gamma_{\zeta_i \oplus e_{p_{j'}^i}} \wedge \Gamma_{\eta \oplus e_{p_{j'}^i}}$ are linearly independent. For i, i' such that $d_i = 2$ and $d_{i'} > 2$, $2\Gamma_{\zeta_i \oplus e_{p_1^i}} \wedge \Gamma_{\eta \oplus e_{p_1^i}}$ and any $\Gamma_{\zeta_{i'} \oplus e_{p_j^{i'}}} \wedge \Gamma_{\eta \oplus e_{p_j^{i'}}}$ are linearly independent as $\zeta_i \oplus \eta \neq \zeta_{i'} \oplus \eta$, contradicting the assumption that $\{\zeta_1, ..., \zeta_m\}$ are distinct. Similarly, for i, i' such that $d_i > 2$ and $d_{i'} > 2$, any $\Gamma_{\zeta_i \oplus e_{p_j^i}} \wedge \Gamma_{\eta \oplus e_{p_j^i}}$ and any $\Gamma_{\zeta_{i'} \oplus e_{p_{j'}^{i'}}} \wedge \Gamma_{\eta \oplus e_{p_{j'}^{i'}}}$ are linearly independent as $\zeta_i \oplus \eta \neq \zeta_{i'} \oplus \eta$.

We conclude that (10), after consolidating sums for which $d_i = 2$, consists of a nonzero number of linearly independent wedge products of columns indexed by S, so (10) is indeed a nontrivial linear relation among the wedge products $\Gamma_s \wedge \Gamma_{s'}$ for $s, s' \in S$. But all columns indexed by S are linearly independent by assumption, so this linear relation contradicts Lemma 3.8 and the linear independence of columns indexed by S.

For the second part of Lemma 3.9, we claim that (10) is still a nontrivial relation. Pick any $k \neq j^*$ inside $[d_0]$. Because $d_0 \geq 4$,

$$\zeta_0 \oplus e_{p_k^0}, \eta \oplus e_{p_k^0} \neq \eta \oplus e_{p_{j^*}^0}. \tag{11}$$

We have already taken care of the case where $\Gamma_{\zeta_0 + e_{p_{j^*}^0}} \notin \mathrm{span}(S')$ above, so suppose instead that

$$\Gamma_{\zeta_0 + e_{p_{j^*}^0}} = \sum_{s \in S'} b_s \Gamma_s. \tag{12}$$

If we consolidate sums for which $d_i = 2$ in (10) as above and substitute (12) into the resulting equation, the wedge products that (10) now contains also include ones of the form $\Gamma_s \wedge \Gamma_{\eta + e_{p_{j^*}^0}}$, which cannot be linearly dependent with $\Gamma_{\zeta_0 \oplus e_{p_k^0}} \wedge \Gamma_{\eta \oplus e_{p_k^0}}$ by (11).

Every other wedge product in (10) is of the form $\Gamma_{\zeta_i \oplus e_{p_j^i}} \wedge \Gamma_{\eta \oplus e_{p_j^i}}$ and also cannot be linearly dependent with $\pm\Gamma_{\zeta_0 \oplus e_{p_k^0}} \wedge \Gamma_{\eta \oplus e_{p_k^0}}$ or else, as before, we'd find that $\zeta_i \oplus \eta = \zeta_k \oplus \eta$, contradicting the assumption that $\{\zeta_1, ..., \zeta_m\}$ are distinct.

It follows that if we rewrite the left-hand side of (10) in the form $\sum_{s,s' \in S'} b_{s,s'} \Gamma_s \wedge \Gamma_{s'}$ (uniquely because by Lemma 3.8 the $\Gamma_s \wedge \Gamma_{s'}$ are linearly independent), $b_{\zeta_0 \oplus e_{p_k^0}, \eta \oplus e_{p_k^0}} \neq 0$. So (10) is still a nontrivial linear relation, contradicting Lemma 3.8 and the linear independence of columns indexed by S. □

4 Rigidity and Cluster Existence

In this section, we will prove the rigidity theorem and the cluster existence theorem, informally stated as Theorems 1.2 and 1.3. Now that we have introduced the appropriate terminology, we

first state both precisely.

Theorem 4.1 (Rank Rigidity - formal)**.** *If* Γ *is a* $2^\ell \times 2^m$ *pseudo-signature, then* $\text{rank}(\Gamma)$ *is a power of 2. Equivalently, for all* $\kappa \geq 1$*,*

$$\text{rank}(\Gamma) \geq 2^\kappa + 1 \Rightarrow \text{rank}(\Gamma) \geq 2^{\kappa+1} \tag{13}$$

Theorem 4.2 (Cluster Existence - formal)**.** *If* $\Gamma = \underline{G}(t)$ *for some matchgate* G *of arity* $n\ell$*, and* $\text{rank}(\Gamma) \geq k$*, then there is a* $\lceil \log_2 k \rceil$*-cluster submatrix of full rank.*

To prove Theorem 4.2, we claim it is enough to show the following:

Theorem 4.3. *If* Γ *is a* $2^\ell \times 2^m$ *pseudo-signature such that* $\text{rank}(\Gamma) \geq k$*, then there exists a* $(\lceil \log_2 k \rceil, m)$*-cluster* Z *for which* Γ_Z *is of full rank.*

Equivalently, for all $\kappa \geq 1$*,*

$$\text{rank}(\Gamma) \geq 2^\kappa + 1 \Rightarrow \exists\ (\kappa+1, m)\text{-cluster } Z \text{ such that } \text{rank}(\Gamma_Z) = 2^{\kappa+1} \tag{14}$$

Proof of Theorem 4.2. Apply Theorem 4.3 to Γ to obtain Γ_Z. By Observation 1, Γ_Z^T is also a pseudo-signature. Apply Theorem 4.3 to Γ^T to get the desired cluster submatrix. □

Note that for fixed κ, (14) implies (13). We will jointly prove Theorem 4.1 and Theorem 4.3 by induction on k. Cai and Fu have already shown both for $k \leq 4$; we take these results as our base case. We complete the following two inductive steps.

Inductive Step 1. *If implication* (14) *holds for* $1 \leq \kappa \leq K-1$*, then implication* (13) *holds for* $1 \leq \kappa \leq K$*.*

Inductive Step 2. *If implication* (13) *holds for* $1 \leq \kappa \leq K$ *and implication* (14) *holds for* $\kappa \leq K-1$*, then implication* (14) *also holds for* $\kappa = K$*.*

Note that once we have proven the rigidity and cluster existence theorems, we additionally obtain the following.

Corollary 4.4. *If* G *is a full-rank matchgate signature on domain size* k*, it is only realizable over bases* M *of rank at most* $2^{\lfloor \log_2 k \rfloor}$*.*

Proof. If k is a power of two, the claim is trivial. If not, assume G is a generator (standard signatures of recognizer are also pseudo-signatures, so the argument in that case is analogous). If $\text{rank}(G(t)) = k$ and $\text{rank}(M) \geq 2^{\lfloor \log_2 k \rfloor} + 1$, then by Lemma 2.14, $2^{\lfloor \log_2 k \rfloor} + 1 \leq \text{rank}(\underline{G}(t)) \leq k$. But Theorem 4.1 would then imply $\text{rank}(\underline{G}(t)) \geq 2^{\lfloor \log_2 k \rfloor + 1} > k$, a contradiction. □

4.1 Rank Rigidity Theorem

In this subsection, we complete the former inductive step, and in the next, we complete the latter.

Before we prove Inductive Step 1 in its entirety, we take care of the case where $m = K+1$. While this might appear to be the simplest case because m is minimal, it will turn out that cases where m is greater will reduce to this case. For this reason, the wedge product machinery introduced in Section 3 is used exclusively in the proof of this case.

Once we take care of this case, we will essentially show that if standard signature Γ is any wider, i.e. if $m > K+1$, then if Z is a cluster of size greater than 2^K indexing columns of rank at least 2^K+1, we can always find a proper subcluster also indexing columns of rank at least 2^K+1, or else the matchgate identities would erroneously imply that certain columns which are known to be linearly independent are linearly dependent.

We begin with the case of $m = K+1$.

Theorem 4.5. *If Γ is a $2^\ell \times 2^{K+1}$ pseudo-signature such that $\operatorname{rank}(\Gamma) \geq 2^K+1$, then $\operatorname{rank}(\Gamma) = 2^{K+1}$.*

Proof. Because we assume implication (14) holds for $\kappa = K-1$, Γ contains a (K,m)-cluster $Z = s \oplus \{e_1, ..., \hat{e_j}, ..., e_{K+1}\}$ of linearly independent columns; denote the even indices of Z by Z_0 and the odd ones by Z_1. Because $\operatorname{rank}(\Gamma) > 2^K$, there exists $t \notin Z$ for which $\Gamma_t \notin \operatorname{span}(Z)$. Denote the parity of t by $b \in \{0,1\}$, and denote by $\bar{b}$ the opposite parity.

Select any $t' = t \oplus e_{i^*}$ for $i^* \neq j$ and apply Lemma 3.9 to $\zeta_0 = t'$, $\eta = t \oplus e_j$, $T = Z_{\bar{b}}$ to conclude that $\Gamma_{t'} \notin \operatorname{span}(Z_{\bar{b}})$.

Let S_{d_1,d_2} denote the set of column indices $\zeta \notin Z$ for which $\operatorname{wt}(\zeta \oplus t) = d_1$ and $\operatorname{wt}(\zeta \oplus t') = d_2$. Note that because $\operatorname{wt}(t \oplus t') = 1$, S_{d_1,d_2} is empty if $|d_1 - d_2| \neq 1$.

We will show by induction on D that the columns indexed by $\left(\bigcup_{d=0}^{D} S_{d,d+1} \cup S_{d+1,d}\right) \cup Z$ are linearly independent for all $0 \leq D \leq K$. The definition of t and the argument above for t' give the base case of $D = 0$.

For the inductive step, for each d let d_0 and d_1, respectively denote the even and odd value in $\{d, d+1\}$. As columns indexed by S_{d_0,d_1} and columns indexed by S_{d_1,d_0} have opposite parity, it suffices to show that the columns indexed by $T_D^0 := \left(\bigcup_{d=0}^{D} S_{d_0,d_1}\right) \cup Z_b$ are linearly independent, and that the columns indexed by $T_D^1 := \left(\bigcup_{d=0}^{D} S_{d_1,d_0}\right) \cup Z_{\bar{b}}$ are linearly independent.

Within this inductive step, we will further induct on the elements within S_{D_0,D_1} and S_{D_1,D_0}. Specifically, suppose we have already proven that for some subset $S'_{D_0,D_1} \subset S_{D_0,D_1}$, all columns indexed by $T_{D-1}^0 \cup S'_{D_0,D_1}$ are linearly independent, and that for $S'_{D_1,D_0} := \{u \oplus e_{i^*} : u \in S'_{D_0,D_1}\} \subset S_{D_1,D_0}$, all columns indexed by $T_{D-1}^1 \cup S'_{D_1,D_0}$ are linearly independent. Select any $u \notin S'_{D_0,D_1}$ and apply Lemma 3.9 to $\zeta_0 = u$, $\eta = t' \oplus e_j$, $T = T_{D-1}^0 \cup S'_{D_0,D_1}$ to see that $\Gamma_u \notin \operatorname{span}(T_{D-1}^0 \cup S'_{D_0,D_1})$. Note that when $\operatorname{wt}(\zeta_0 \oplus \eta) \geq 4$, we do not yet know that $\Gamma_{\zeta_0 \oplus e_{i^*}} = \Gamma_{u \oplus e_{i^*}}$ lies outside $\operatorname{span}(T_{D-1}^1 \cup S'_{D_1,D_0})$, that is, we do not know whether all the columns indexed by the set S defined in Lemma 3.9 are linearly independent, but the second part of Lemma 3.9 says that we may still conclude that $\Gamma_u \notin \operatorname{span}(T_{D-1}^0 \cup S'_{D_0,D_1})$ because the columns indexed by $S \backslash \{u \oplus e_{i^*}\}$ are linearly independent.

Lastly, apply Lemma 3.9 to $\zeta_0 = u \oplus e_{i^*}$, $\eta = t \oplus e_j$, $T = T_{D-1}^1 \cup S'_{D_1,D_0}$ to see that $\Gamma_{u \oplus e_{i^*}} \notin \operatorname{span}(T_{D-1}^1 \cup S'_{D_1,D_0})$. Note that here we only need to invoke the first part of Lemma 3.9 we already know that $\Gamma_{\zeta_0 \oplus e^*} = \Gamma_u$ lies outside $\operatorname{span}(T_{D-1}^0 \cup S'_{D_0,D_1})$. □

We are now ready to complete Inductive Step 1.

Completion of Inductive Step 1. As we remarked earlier, implication (14) for a fixed value of κ implies implication (13) for that value of κ, so we just need to show that implication (13) also holds for $\kappa = K$.

Suppose to the contrary that there exists pseudo-signature Γ of rank k such that $2^K + 1 \leq k < 2^{K+1}$. Without loss of generality, we may assume that for all clusters $Z \subsetneq \{0,1\}^m$, $\operatorname{rank}(\Gamma_Z) \leq 2^K$; otherwise, replace Γ by Γ_Z for some small enough cluster Z such that $\operatorname{rank}(\Gamma_Z) \geq 2^K + 1$ and $\operatorname{rank}(\Gamma_{Z'}) \leq 2^K$ for all subclusters $Z' \subsetneq Z$. Furthermore, by Theorem 4.5, we may assume $m > K+1$.

Lemma 4.6. *If $Z = s + \{e_{p_1}, ..., e_{p_K}\}$ is a (K, m)-cluster of linearly independent columns in Γ, then any column Γ_t for which $t_i = s_i$ for some $i \neq p_1, ..., p_K$ lies in the span of the columns indexed by Z.*

Proof. If to the contrary there existed such a Γ_t not lying in the span of Z so that $t_i = s_i$ for some $i \neq p_1, ..., p_K$, then if Z' is the $(m-1, m)$-cluster of column indices ζ for which $\zeta_i = s_i$, $\Gamma_{Z'}$ contains t and all of Z and thus has rank at least $2^K + 1$, contradicting our assumption on the ranks of the proper subclusters of Γ. □

By the inductive hypothesis, Γ contains a (K, m)-cluster Z of linearly independent columns $s + \{e_{p_1}, ..., e_{p_K}\}$. As s is only uniquely defined modulo the bits in positions $p_1, ..., p_K$, we will leave those bits of s unspecified for now.

Because $k > 2^K$, there exists $t \notin Z$ for which all columns indexed by $Z \cup \{t\}$ are linearly independent. Moreover, by Lemma 4.6, $t_j = \overline{s_j}$ for all $j \notin \{p_1, ..., p_K\}$. Let Z' denote the cluster $t + \{e_{p_1}, ..., e_{p_K}\}$. Set $s_i = t_i$ for $i \in \{p_2, ..., p_K\}$; we will set s_{p_1} to be 0 or 1 depending on the parity of the number $m - K$ of bits outside of positions $p_1, ..., p_K$.

Case 1. *$m - K$ is even.*

Set $s_{p_1} = t_{p_1}$ so that s and t have the same parity.

Claim 4.7. *If $u \notin Z'$ and $u_i = s_i$ for each $i \in \{p_1, ..., p_K\}$, then Γ_u and Γ_s are linearly dependent.*

Proof. For each $i \in \{p_1, ..., p_K\}$ and $j \notin \{p_1, ..., p_K\}$, let T_i denote the cluster of all column indices u for which $u_i = s_i$, and let T_i^j denote the cluster of all column indices u for which $u_i = s_i$ and $u_j = s_j$. Let $Z_i = Z \cap T_i$; obviously $Z_i \subset T_i^j \subset T_i$.

Because T_i^j is a cluster properly contained in $\{0, 1\}^m$, we inductively know that $\text{rank}(\Gamma_{T_i^j}) \leq 2^K$. And because $Z_i \subset T_i^j$, $\text{rank}(\Gamma_{T_i^j}) \geq 2^{K-1}$. But if $\text{rank}(\Gamma_{T_i^j}) \geq 2^{K-1} + 1$, then by inductive hypothesis (14) applied to $\Gamma_{T_i^j}$ for $\kappa = K - 1$, $\text{rank}(\Gamma_{T_i^j}) \geq 2^K$. In other words, $\text{rank}(\Gamma_{T_i^j})$ is either 2^{K-1} or 2^K.

We will show that the latter is impossible. Suppose to the contrary that $\text{rank}(\Gamma_{T_i^j}) = 2^K$.

Then because $T_i^j \subset T_i$ and $\text{rank}(\Gamma_{T_i}) = 2^K = \text{rank}(\Gamma_{T_i^j})$, it follows that $\text{span}(T_i) = \text{span}(T_i^j)$. For any $u \in T_i^j$, $\Gamma_u \in \text{span}(Z)$ by Lemma 4.6, so

$$\text{span}(Z) \supset \text{span}(T_i^j) = \text{span}(T_i).$$

But T_i contains t, and by definition $\Gamma_t \notin \text{span}(Z)$, a contradiction.

We conclude that $\text{rank}(\Gamma_{T_i^j}) = 2^{K-1}$. Then because $Z_i \subset T_i^j$ and $\text{rank}(\Gamma_{Z_i}) = 2^{K-1} = \text{rank}(\Gamma_{T_i^j})$, it follows that $\text{span}(Z_i) = \text{span}(T_i^j)$.

In particular, all columns indexed by $\bigcap_{k=1}^K T_{p_k}^j$ lie in $\bigcap_{k=1}^K \text{span}(Z_{p_k}) = \text{span}(\{s\})$. Our choice of j was arbitrary, so we get the desired claim. □

From the above claim and the fact that we're assuming $m > K + 1$, we conclude that $\Gamma_{s \oplus e_j}$ for any $j \notin \{p_1, ..., p_K\}$ lies in the span of Γ_s. But s and $s \oplus e_j$ are of opposite parity, so by the parity condition, $\Gamma_{s \oplus e_j} = 0$ for all such j. Applying Corollary 3.6 to s and t, it follows that Γ_s and Γ_t are linearly dependent, a contradiction.

Case 2. *$m - K$ is odd.*

Set $s_{p_1} = \overline{t_{p_1}}$ so that s and t have the same parity.

Claim 4.8. *For any $u \in \{0,1\}^m$, if $u \notin Z'$ and $u_i = s_i$ for all $i \in \{p_2, ..., p_K\}$, then:*

1. *If u and s have the same parity, then Γ_u and Γ_s are linearly dependent.*
2. *If u and s have the opposite parity, then Γ_u and $\Gamma_{s \oplus e_{p_1}}$ are linearly dependent.*

Proof. The proof is the same as that of Claim 4.7, the only subtlety being that s and t now only necessarily agree on bits $p_2, ..., p_K$. By the argument there, all columns indexed by T_i^j lie in $\text{span}(Z_i)$ for $i = p_2, ..., p_K$. In particular, for all $j \notin \{p_1, ..., p_K\}$, all columns indexed by $\bigcap_{k=2}^K S_{p_k}^j$ lie in $\bigcap_{k=2}^K \text{span}(Z_{p_k}) = \text{span}(\{s, s \oplus e_{p_1}\})$.

So given $u \notin Z'$, write $\Gamma_u = \alpha\Gamma_s + \beta\Gamma_{s \oplus e_{p_1}}$. If u and s have the same parity, $\beta = 0$ by the parity condition, so $\Gamma_u \in \text{span}(\{s\})$. If u and s have the opposite parity, $\alpha = 0$ by the parity condition, so $\Gamma_u \in \text{span}(\{s \oplus e_{p_1}\})$. □

Pick any $j \notin \{p_1, ..., p_K\}$ and define $s^* = s \oplus e_j$ and $t^* = t \oplus e_j$. s^* and t^* both satisfy the hypotheses of Claim 4.8 and have parity opposite to that of s, so by the latter case of Claim 4.8, they are both linearly dependent with $\Gamma_{s \oplus e_{p_1}}$. But $\Gamma_{s \oplus e_{p_1}} \neq 0$ because $s \oplus e_{p_1} \in Z$ and the columns indexed by Z are linearly independent, so Γ_{s^*} and Γ_{t^*} are linearly dependent with each other.

To show Γ_s and Γ_t are linearly dependent, we wish to apply Corollary 3.7 to s^*, t^*, noting that $s^* \oplus t^* = e_j \oplus \sum_{j' \notin \{j, p_2, ..., p_k\}} e_j$.

For any $j' \notin \{j, p_2, ..., p_K\}$, note that $s^* \oplus e_{j'}$ and $t^* \oplus e_{j'}$ both satisfy the hypotheses of Claim 4.8 and have the same parity as s, so by the former case of Claim 4.8, they are both linearly dependent with Γ_s. But $\Gamma_s \neq 0$ because $s \in Z$ and the columns indexed by Z are linearly independent, so $\Gamma_{s^* \oplus e_{j'}}$ and $\Gamma_{t^* \oplus e_{j'}}$ are linearly dependent with each other.

Applying Corollary 3.7 to s^* and t^*, it follows that Γ_s and Γ_t are linearly dependent, a contradiction. □

4.2 Existence of Cluster Submatrix

Completion of Inductive Step 2. As in inductive step 1, we may assume without loss of generality that for all clusters $Z \subsetneq \{0,1\}^m$, $\text{rank}(\Gamma_Z) \leq 2^K$. If $m = K + 1$, then by the inductive hypothesis that (13) holds for $\kappa = K$, we're done. So suppose $m > K + 1$.

By the second part of the inductive hypothesis, implication (14) holds for $1 \leq \kappa \leq K - 1$, so Γ contains a (K, m)-cluster Z of linearly independent columns $s + \{e_{p_1}, ..., e_{p_K}\}$.

As in inductive step 1, we can apply Lemma 4.6 to Z to see that all columns outside the span of the columns indexed by Z must be indexed by $Z' = t + \{e_{p_1}, ..., e_{p_K}\}$, where $t = \left(s \oplus \bigoplus_{i \neq p_1, ..., p_K} e_i\right)$. But $|Z'| = |Z| = 2^K$, and $\text{rank}(\Gamma) \geq 2^{K+1}$ by implication (13) for $\kappa = K$, so the columns indexed by $Z \cup Z'$ are linearly independent. Because $m > K + 1$, there exist columns not indexed by either Z or Z', and by Lemma 4.6 applied once to Z and once to Z', these columns are in both $\text{span}(Z)$ and $\text{span}(Z')$ and thus must be zero.

If s and t are of the same parity, apply Corollary 3.6 to s and t to find that Γ_s and Γ_t are linearly dependent, a contradiction.

If s and t are of opposite parity, apply Corollary 3.7 to $s \oplus e_j$ and $t \oplus e_j \oplus e_{p_1}$ for any $j \notin \{p_1, ..., p_K\}$ to find that Γ_s and $\Gamma_{t \oplus e_{p_1}}$ are linearly dependent, a contradiction. □

5 Group Property of Standard Signatures

We will now prove the following generalization of the group property result over domain size 4 due to Cai and Fu (Theorem 5.5, [5]):

Theorem 5.1. *If G is a generator matchgate of arity Kn with standard signature $\underline{G}$, and* $\mathrm{rank}(\underline{G}(t)) = 2^K$ *for some t, then there exists a recognizer matchgate of arity Kn such that $\underline{G}(t)\underline{R}(t) = I_{2^K}$.*

Roughly, we invoke Theorem 4.3 to obtain a full-rank K-cluster submatrix G' of $\underline{G}(t)$ with column indices belonging to cluster $\zeta + \{e_{p_1}, ..., e_{p_K}\}$. Assume without loss of generality that $\zeta_{p_i} = 0$ for all $i \in [K]$. We will show that the matrix obtained by replacing G' in $\underline{G}(t)$ with $(G')^{-1}$ and the remaining entries with zeroes is the standard signature of some arity-Kn recognizer.

We first fix some notation. Denote $\underline{G}(t)$ by Γ. Suppose that nodes $p_1 < \cdots < p_m \in [Kn]$ belong to blocks before the t-th, and nodes $p_{m+1} < \cdots < p_K \in [Kn]$ belong to blocks after the t-th. For expository purposes, we wish to use a particular permutation $(q_1, ..., q_K)$ of $(p_1, ..., p_K)$, so for $i \leq m$, let $q_i = p_{m-i+1}$, and for $i > m$, let $q_i = p_{K+m-i+1}$ (see Figure 1). If the column indices of Γ are of the form $i_1 \cdots i_{K(n-1)}$, those of G' are of the form $i_{p_1} \cdots i_{p_K}$.

In [15], Li and Xia gave a constructive proof that in the character theory of matchgates, the $2^K \times 2^K$ *character* matrices of invertible K-input, K-output matchgates form a group under matrix multiplication. One can check that their construction carries over to show that the $2^K \times 2^K$ standard signatures of such matchgates likewise form a group, but unfortunately this is not enough to prove Theorem 5.1, as G' alone is merely a pseudo-signature and may not be realizable as the standard signature of a K-input, K-output matchgate. That said, Theorem 5.1 can still be proved with minor modifications to Li and Xia's approach.

We begin with a toy example motivating the notation in the previous paragraphs. Suppose that for each $i \in [K]$, there exists an edge of weight 1 such that the i-th external node in block t and external node q_i are both incident only to this edge. Note that in this case, G' is a symmetric permutation matrix and thus equal to its own inverse. We can easily construct a recognizer R out of G for which $\underline{G}(t)\underline{R}(t) = I_{2^K}$ as follows. Remove all non-external nodes of G, as well as all edges incident to nodes outside of block t and nodes $q_1, ..., q_K$. For external node i outside of block t such that $i \notin \{q_1, ..., q_K\}$, if $\zeta_i = 0$, attach a distinct edge of weight 1 to node i and designate the other endpoint of the edge as the ith external node of R; if $\zeta_i = 1$, attach a distinct path graph of length 2 consisting of two edges of weight 1, and denote the other endpoint of the path as the ith external node of R. By construction, in the $2^{K(n-1)} \times 2^K$ matrix $\underline{R}(t)$, the submatrix indexed by rows $q_1, ..., q_K$ is equal to G', and all other entries are zero. Because $G' = (G')^{-1}$, $\underline{G}(t)\underline{R}(t) = I_{2^K}$ as desired.

See Figure 1 for an example of this construction.

Definition 5.2. For $i \in [K]$, if $i \leq m$ (resp. $i > m$), G is *reduced at i* if there exists an edge of weight 1 in G connecting the i-th external node in block t and external node q_i such that these two nodes are both incident only to that edge.

To prove Theorem 5.1, it is enough to reduce to the special case of the toy example above where Γ is realizable by a generator G reduced at all $i \in [K]$. The rest of this section will be dedicated to proving the following:

Lemma 5.3. *If Γ is the standard signature of a generator of arity Kn reduced at $1, ..., i$, there exist nonsingular K-input, K-output transducers $L_1, ..., L_r$ and $K(n-1)$-input, $K(n-1)$-output transducers $R_1, ..., R_s$ such that $\underline{L_r} \cdots \underline{L_1} \cdot \Gamma \cdot \underline{R_1} \cdots \underline{R_s}$ is the standard signature of a generator reduced at $1, ..., i+1$.*

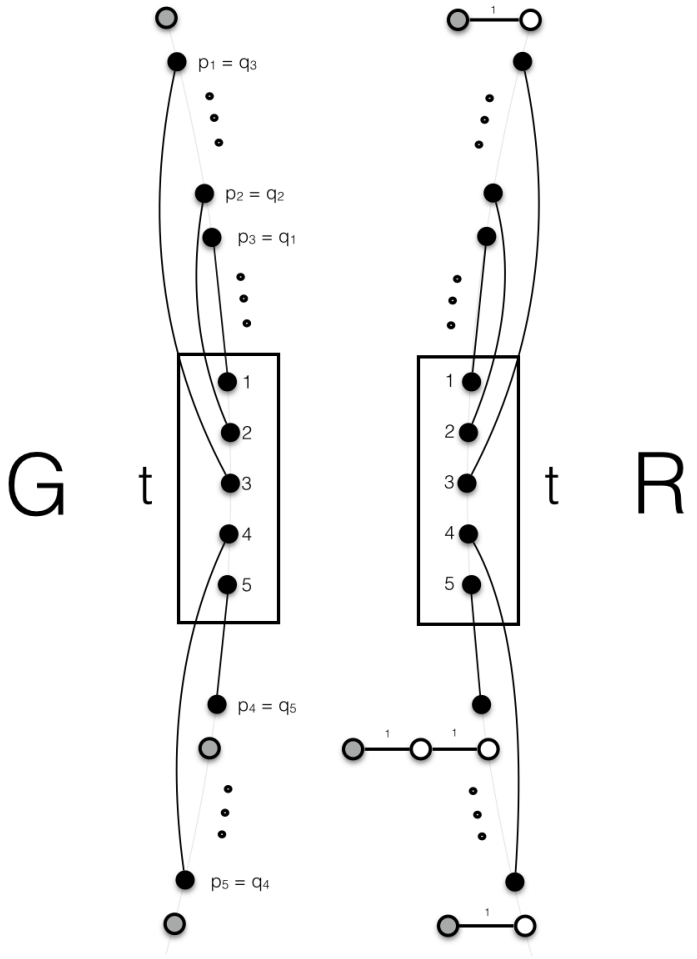

Figure 1: Toy example of G reduced at $i = 1, ..., K$. Here, $\ell = 5$, $K = 5$, $m = 3$. Black nodes in G denote nodes $1, ..., K$ in block t and nodes $q_1, ..., q_K$. External nodes of G and R are shown in black/gray.

We first give a sufficient characterization of standard signatures of matchgates reduced at $1, ..., i$ in terms of the entries of their standard signatures.

The following terminology is borrowed from [15].

Definition 5.4. Let M be any $2^r \times 2^c$ matrix whose rows and columns are indexed by $\sigma \in \{0,1\}^r$ and $\tau \in \{0,1\}^c$ respectively. Γ^σ_τ is an *edge entry of* M iff $r + c - 2 \leq \mathrm{wt}(\sigma) + \mathrm{wt}(\tau) < r + c$.

Lemma 5.5. *Γ is the standard signature of a generator G that is reduced at i if Γ satisfies the following:*

1. $(G')^{1^K}_{1^K} = (G')^{1^K \oplus e_i}_{1^K \oplus e_{q_i}} = 1$.
2. $(G')^\sigma_\tau = 0$ *for all other edge entries of G' for which* $\sigma \in \{1^K, 1^K \oplus e_i\}$ *or* $\tau \in \{1^K, 1^K \oplus e_{q_i}\}$.

To show this, it suffices to prove the following useful consequence of the matchgate identities, first observed in [4] and translated below to our setting of standard signatures in matrix form.

Lemma 5.6 (Theorem 4.2, [4])**.** *If $(G')^{1^K}_{1^K} \neq 0$, the entries of G' are uniquely determined by $(G')^{1^K}_{1^K}$ and the edge entries of G'.*

Proof. Assume that these entries uniquely determine all entries Γ^σ_τ for which $\mathrm{wt}(\sigma) + \mathrm{wt}(\tau) \geq m$ for some $m \leq n-2$. We proceed by downward induction on m (by the parity condition, if m is even, the case of $m+1$ follows immediately from that of m). Then for $\sigma, \tau \in \{0,1\}^K$ such that $\mathrm{wt}(\sigma) + \mathrm{wt}(\tau) = m-2$, apply (2) from Theorem 3.1 to $\sigma := \sigma$, $\zeta := \tau$, $\tau := 1^K$, and $\eta := 1^K$. One can check that the resulting identity consists of $(G')^\sigma_\tau \cdot (G')^{1^K}_{1^K}$ and terms which have already been uniquely determined by the inductive hypothesis, so because $(G')^{1^K}_{1^K} \neq 0$, we conclude that $(G')^\sigma_\tau$ is also uniquely determined. □

Observation 2. *If Γ is the standard signature of a generator reduced at i, $(G')^{\sigma \oplus e_i}_{\tau \oplus e_{q_i}} = (G')^\sigma_\tau$, and if $\sigma_i \neq \tau_{q_i}$, $(G')^\sigma_\tau = 0$.*

Proof. By hypothesis, external node i of block t and external node q_i are only connected to each other. If $\sigma_i \neq \tau_{q_i}$, $(G')^\sigma_\tau$ counts the number of perfect matchings of Γ where, among other conditions, exactly one of these two nodes is removed, and no such matching exists. On the other hand, if $\sigma_i = \tau_{q_i}$, then $(G')^\sigma_\tau$ and $(G')^{\sigma \oplus e_i}_{\tau \oplus e_{q_i}}$ count the number of perfect matchings in which, among other conditions, both of these two nodes are removed, or neither is. The number of perfect matchings in either scenario is the same. □

Let G'_i be the $2^{K-i} \times 2^{K-i}$ submatrix of G' consisting of entries $(G')^\sigma_\tau$ for which $\sigma_j = 0$ for $j = 1, ..., i$ and $\tau_j = 0$ for $j = q_1, ..., q_i$. If the row and column indices of G' are of the form $i_1 \cdots i_K$ and $i_{p_1} \cdots i_{p_K}$ respectively, the row and column indices of G'_i are of the form $i_{i+1} \cdots i_K$ and $i_{p_{i+1}} \cdots i_{p_K}$ respectively. When referring to the row (resp. column) of G' containing a row $i_{i+1} \cdots i_K$ (resp. column $i_{p_{i+1}} \cdots i_{p_K}$) of G'_i, we use the notation $0^i \circ i_{i+1} \cdots i_K$ (resp. $0^i \circ i_{p_{i+1}} \cdots i_{p_K}$) to denote its index in G'. For example, column $0^i \circ 1^{K-i}$ of G' is the column of G' indexed by $\sigma \in \{0,1\}^K$ for which $\sigma_{q_1} = \cdots = \sigma_{q_i} = 0$ and $\sigma_{q_{i+1}} = \cdots = \sigma_{q_K} = 1$, and this contains column 1^{K-i} of G'_i.

Corollary 5.7. *If Γ is the standard signature of a generator that is reduced at $1, ..., i$, then Γ is the standard signature of a generator reduced at $1, ..., i+1$ if Γ satisfies the following:*

1. $(G'_i)^{1^{K-i}}_{1^{K-i}} = (G'_i)^{1^{K-i} \oplus e_{i+1}}_{1^{K-i} \oplus e_{q_{i+1}}} = 1$
2. $(G'_i)^\sigma_\tau = 0$ *for all other edge entries of G'_i such that $\sigma \in \{1^{K-i}, 1^{K-i} \oplus e_{i+1}\}$ or $\tau \in \{1^{K-i}, 1^{K-i} \oplus e_{q_{i+1}}\}$.*

Proof. Apply Observation 2 to each of $1, ..., i$, and invoke Lemma 5.5. □

Proof of Lemma 5.3. We execute the transformation $\Gamma = \Gamma^{(0)} \Rightarrow \Gamma^{(1)} \Rightarrow \Gamma^{(2)} \Rightarrow \Gamma^{(3)} \Rightarrow \Gamma^{(4)}$ outlined below.

1. $(\Gamma^{(0)} \Rightarrow \Gamma^{(1)})$: Turn the entry indexed by $(1^{K-i}, 1^{K-i})$ in G'_i to 1.
2. $(\Gamma^{(1)} \Rightarrow \Gamma^{(2)})$: Turn edge entries of G'_i in row or column 1^{K-i} to 0.
3. $(\Gamma^{(2)} \Rightarrow \Gamma^{(3)})$: Turn entry $(1^{K-i} \oplus e_{i+1}, 1^{K-i} \oplus e_{q_{i+1}})$ in G'_i to 1.

4. ($\Gamma^{(3)} \Rightarrow \Gamma^{(4)}$): Turn all other edge entries in G'_i in row $1^{K-i} \oplus e_{i+1}$ or column $1^{K-i} \oplus e_{q_{i+1}}$ to zero.

We need not care what these transformations do to entries outside of G', but we must ensure they preserve the fact that Γ is the standard signature of a generator reduced at $1, ..., i$. To do this, for each matrix M by which we left- or right-multiply Γ, if σ does not index a row (resp. column) of G'_i, the only nonzero entry of M in row (resp. column) σ will be 1 in column (resp. row) σ.

In each step j, we will for convenience refer to $\Gamma^{(j-1)}$ as Γ.

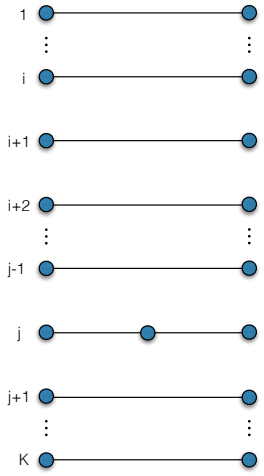

(a) C_j

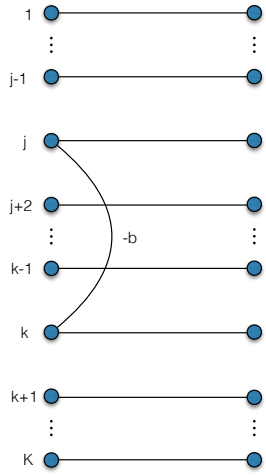

(b) $L_3^{j,k}$

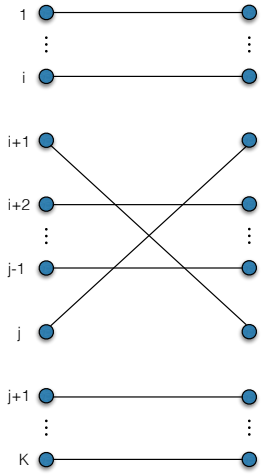

(c) L_4

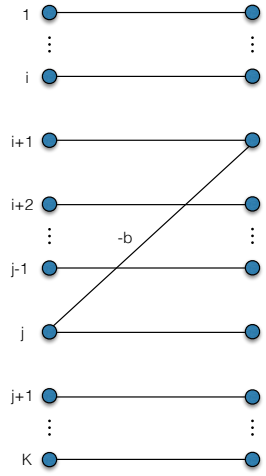

(d) L_5^j

Figure 2: Transducers realizing row/column operations in steps 1-4

Step 1. *($\Gamma^{(0)} \Rightarrow \Gamma^{(1)}$): Turn the entry indexed by $(1^{K-i}, 1^{K-i})$ in G'_i to 1.*

We first show how to move a nonzero entry $c := (G'_i)^{\sigma^*}_{\tau^*}$ of G'_i into entry $(1^{K-i}, 1^{K-i})$ of G'_i.

For each j for which $i < j \leq K$, we would like a $2^K \times 2^K$ standard signature C_j such that left-multiplication of Γ by C_j interchanges row σ in G' with $\sigma \oplus e_j$ for all $\sigma \in \{0,1\}^K$, and a $2^{K(n-1)} \times 2^{K(n-1)}$ standard signature D_j such that right-multiplication of Γ by D_j interchanges column τ in G' with $\tau \oplus e_{q_j}$. We could then define $L_1 = \prod_{j:\sigma_j^*=0} C_j$ and $R_1 = \prod_{j:\tau_j^*=0} D_j$, and $L_1 \cdot \Gamma R_1$ would have nonzero entry c at index $(1^{K-i}, 1^{K-i})$ of G_i' and still be the standard signature of a matchgate reduced at $1, ..., i$.

C_j (resp. D_j) is the permutation matrix whose only nonzero entry in row $\sigma \in \{0,1\}^K$ (resp. $\sigma \in \{0,1\}^{K(n-1)}$) is 1 in column $\sigma \oplus e_j$ (resp. column $\sigma \oplus e_{q_j}$) if G_i' contains entries from Γ^σ (resp. Γ_σ), and 1 in column σ otherwise. C_j and D_j are certainly nonsingular.

To construct the K-input, K-output transducer realizing C_j as a standard signature, begin with a (K,K)-bipartite graph where for every $\nu \neq j$, left node ν and right node ν are connected by an edge of weight 1. Add an extra vertex between left node j and right node j, and draw a path of length two connecting these three vertices, where both edges of the path have weight 1. This construction is shown in Figure 2a.

The $K(n-1)$-input, $K(n-1)$-output transducer realizing D_j as a standard signature is similarly constructed, the only difference being that the bipartite graph has left and right vertex sets of size $K(n-1)$, and the path of length two is drawn between left node q_j and right node q_j.

Next, we want to scale all of the entries of $L_1 \Gamma R_1$ by a factor of $1/c$, so define $L_2(c)$ to be the diagonal matrix whose entry at index 1^{K-i} is $1/c$ and whose entries at all other indices are 1. Obviously $L_2(c)$ is nonsingular and satisfies the matchgate identities (2). We take $\Gamma^{(1)} = L_2(c) L_1 \Gamma^{(0)} R_1$.

Step 2. *($\Gamma^{(1)} \Rightarrow \Gamma^{(2)}$): Turn edge entries of G_i' in row or column 1^{K-i} to 0.*

We demonstrate how to do this for edge entries in column 1^{K-i}. Firstly, edge entries $(1^{K-i} \oplus e_j, 1^{K-i})$ in G_i' are already zero by the parity condition.[1] To set each of the remaining edge entries in this column to zero, we will proceed in reverse lexicographic order over the rows $1^{K-i} \oplus e_j \oplus e_k$ of G_i' and at each step left-multiply Γ by a matrix $L_3^{j,k}$ which corresponds in Γ to subtracting $b := (G_i')_{1^{K-i}}^{1^{K-i} \oplus e_j \oplus e_k}$ times row $0^i \circ 1^{K-i}$ of G' from row $0^i \circ (1^{K-i} \oplus e_j \oplus e_k)$ of G'.

$L_3^{j,k}$ must be a matrix whose nonzero entries include diagonal entries equal to 1 and entry $(0^i \circ (1^{K-i} \oplus e_j \oplus e_k), 0^i \circ 1^{K-i})$ equal to $-b$. A standard signature satisfying these conditions can be realized by the following matchgate: construct a (K,K)-bipartite graph in which left node ν and right node ν are connected by an edge of weight 1 for all ν, and left nodes j and k are also connected by an edge of weight $-b$. The standard signature of this matchgate is nonsingular. The construction is shown in Figure 2b.

$L_3^{j,k}$ additionally has nonzero entries $(\sigma \oplus e_j \oplus e_k, \sigma)$ equal to $-b$ for all $\sigma \in \{0,1\}^K$, i.e. left-multiplication by $L_3^{j,k}$ corresponds to subtracting b times row σ from row $\sigma \oplus e_j \oplus e_k$. These extraneous side effects do not however affect any of the progress we've made as the rows $1^{K-i} \oplus e_j \oplus e_k$ of G_i' are taken in reverse lexicographic order.

The only issue is that the matchgate we have constructed is not necessarily planar. But by Lemma A.1 in Appendix A, there exists a planar matchgate with standard signature equal to $L_3^{j,k}$, except at nonzero off-diagonal entries other than $(0^i \circ (1^{K-i} \oplus e_j \oplus e_k), 0^i \circ 1^{K-i})$, where it may differ from $L_3^{j,k}$ by a factor of -1, but by the reasoning in the previous paragraph, this does not matter. Denote the standard signature of this planar matchgate by $L_3'^{j,k}$.

[1] As characters of general matchgates with omittable nodes do not satisfy the parity condition necessarily, the proof of the group property in [15] requires an extra construction to turn edge entries $(1^{K-i} \oplus e_j, 1^{K-i})$ to zero. This is an instance where our proof of the group property for signatures is easier than that for characters.

To achieve step 2 for edge entries in *rows* 1^{K-i} as well, we can define matrices $R_3'^{j,k}$ analogously. We can thus take $\Gamma^{(2)} = \left(\prod_{i+2\le j,k\le K} L_3'^{j,k}\right)\cdot\Gamma^{(1)}\cdot\left(\prod_{i+2\le j,k\le K} R_3'^{j,k}\right)$, where the indexing in the products respects the abovementioned reverse lexicographic order.

Step 3. *($\Gamma^{(2)} \Rightarrow \Gamma^{(3)}$): Turn entry $(1^{K-i}\oplus e_{i+1}, 1^{K-i}\oplus e_{q_{i+1}})$ in G_i' to 1.*

Note that $c' := (G_i')_{1^{K-i}\oplus e_{q_k}}^{1^{K-i}\oplus e_j}$ must be nonzero for some $j,k\in[K]$ or else (G_i') is singular. As in Step 1, we will first left-multiply Γ by some L_4 to move this nonzero entry to row $1^{K-i}\oplus e_{i+1}$ and then right-multiply by some R_4 to move it to column $1^{K-i}\oplus e_{q_{i+1}}$. Unfortunately, multiplying by C_j or D_j defined in Step 1 would interfere with the progress made so far in Steps 1 and 2.

Instead, L_4 must be a matrix whose only nonzero entry in row (resp. column) $0^i\circ(1^{K-i}\oplus e_j)$ is 1 in column (resp. row) $0^i\circ(1^{K-i}\oplus e_{i+1})$, and whose only nonzero entry in row (resp. column) $0^i\circ(1^{K-i}\oplus e_{i+1})$ is 1 in column (resp. row) $0^i\circ(1^{K-i}\oplus e_j)$. A signature satisfying these conditions can be realized by the following matchgate: construct a (K,K)-bipartite graph in which left node ν and right node ν are connected by an edge of weight 1 for all $\nu\neq j, i+1$. Connect left node j to right node $i+1$ and left node $i+1$ to right node j with edges of weight 1. The standard signature L_4 of this is nonsingular. The construction is shown in Figure 2c.

L_4 also has nonzero entries $(i_1\cdots i_j\cdots i_{i+1}\cdots i_K, i_1\cdots i_{i+1}\cdots i_j\cdots i_K)$, so left-multiplication by L_4 corresponds to switching row $0^i\circ(i_1\cdots i_{i+1}i_{i+2}\cdots i_j\cdots i_K)$ with row $0^i\circ(i_1\cdots i_j i_{i+2}\cdots i_{i+1}\cdots i_K)$ for all $i_{i+1},...,i_K\in\{0,1\}$. Multiplication by L_4 affects neither the progress we've made on entry $(1^{K-i},1^{K-i})$ of G_i' because all bits in the row and column indices are equal, nor the progress on the edge entries in row 1^{K-i} and column 1^{K-i} of G_i' because these get swapped with each other and were already all zero, keeping them equal to zero.

As before, the only issue is that the matchgate constructed is not planar. But by Lemma A.2 in Appendix A, there exists a planar matchgate whose standard signature agrees with L_4 at every entry up to sign. By the above, that the nonzero entries other than $(0^i\circ 1^{K-i}\oplus e_j, 0^i\circ 1^{K-i}\oplus e_{i+1})$ and $(0^i\circ 1^{K-i}\oplus e_{i+1}, 0^i\circ 1^{K-i}\oplus e_j)$ may be -1 does not matter. Furthermore, if either entry $(0^i\circ 1^{K-i}\oplus e_j, 0^i\circ 1^{K-i}\oplus e_{i+1})$ or $(0^i\circ 1^{K-i}\oplus e_{i+1}, 0^i\circ 1^{K-i}\oplus e_j)$ were -1, at worst we may eventually need to replace c' with $-c'$, but the argument still holds. Denote the standard signature of this planar matchgate by L_4'. We can analogously define R_4'.

Next, we scale entry $(1^{K-i}\oplus e_{i+1}, 1^{K-i}\oplus e_{q_{i+1}})$ of G_i' by a factor of $1/c'$, so we take $\Gamma^{(3)} = L_2(c')L_4'\Gamma^{(2)}R_4'$.

Step 4. *($\Gamma^{(3)} \Rightarrow \Gamma^{(4)}$): Turn all other edge entries in G_i' in row $1^{K-i}\oplus e_{i+1}$ or column $1^{K-i}\oplus e_{q_{i+1}}$ to zero.*

We demonstrate how to do this for edge entries in column $1^{K-i}\oplus e_{q_{i+1}}$ of G_i'. To set each of the edge entries in this column to zero, we will proceed in reverse lexicographic order over the row indices $1^{K-i}\oplus e_j$ of G_i' and at each step left-multiply Γ by a matrix L_5^j which subtracts $b := (G_i')_{1^{K-i}\oplus e_{q_{i+1}}}^{1^{K-i}\oplus e_j}$ times row $0^i\circ 1^{K-i}$ of G' from row $0^i\circ(1^{K-i}\oplus e_j)$ of G'.

L_5^j is a matrix whose nonzero entries include diagonal entries equal to 1 and entry $(0^i\circ(1^{K-i}\oplus e_j), 0^i\circ(1^{K-i}\oplus e_{i+1}))$ equal to $-b$. A signature satisfying these conditions can be realized by the following matchgate: construct a (K,K)-bipartite graph in which left node ν and right node ν are connected by an edge of weight 1 for all ν, and connect left node j to right node $i+1$ by an edge of weight $-b$. The standard signature of this matchgate is nonsingular. The construction is shown in Figure 2d.

L_5^j additionally has nonzero entries $(0^i\circ(\sigma\oplus e_j\oplus e_{i+1}), 0^i\circ\sigma)$ equal to $-b$ for all $\sigma\in\{0,1\}^{K-i}$ such that $\sigma_j = 1$ and $\sigma_{i+1} = 0$, i.e. for all such σ, left-multiplication by L_5^j corresponds to

subtracting b times row $0^i \circ \sigma$ of G' from row $0^i \circ \sigma \oplus e_j \oplus e_{i+1}$ of G'. As before, these extraneous side effects do not affect the progress we've made in this step because the edge entries in column $1^{K-i} \oplus e_{q_{i+1}}$ of G'_i are being taken in reverse lexicographic order. They also certainly do not affect row 1^{K-i}, nor do they affect column 1^{K-i} as 1^{K-i} does not satisfy the above criteria for σ.

Again, the issue is that the matchgate constructed is not planar, but by Lemma A.3 in Appendix A, there exists a planar matchgate with standard signature equal to L_5^j except possibly at the nonzero off-diagonal entries other than $(0^i \circ (1^{K-i} \oplus e_j), 0^i \circ (1^{K-i} \oplus e_{i+1}))$, where it may differ by a factor of -1. Denote the standard signature of this planar matchgate by $L_5'^j$.

To achieve step 4 for edge entries in rows $1^{K-i} \oplus e_{i+1}$ as well, we can define matrices $R_5'^j$ analogously. We can thus take $\Gamma^{(4)} = \left(\prod_{i+2 \le j \le K} L_5'^j\right) \cdot \Gamma^{(3)} \cdot \left(\prod_{i+2 \le j \le K} R_5'^j\right)$, where the products respect the abovementioned reverse lexicographic order. □

6 Reducing to Domain Size 2^K

In this section we use Theorem 4.1 to reduce proving a basis collapse theorem over all domain sizes to proving one over domain sizes 2^K. The result we will prove is the following generalization of the main result in [12] whose strategy we follow.

Theorem 6.1. *Suppose Theorem 1.1 has been proven for domain size r. If recognizer signatures $R_1, ..., R_a$ and generator signatures $G_1, ..., G_b$ on domain size $k > r$ belonging to matchgrid Ω are simultaneously realizable on a $2^\ell \times k$ basis M of rank r and R_1 is of full rank, then there exists a basis M' of size at most $\lfloor \log_2 r \rfloor$ on which they are simultaneously realizable.*

We'll need some preliminaries before we can prove this. Express M as $\begin{pmatrix}\alpha_1 & \alpha_2 & \cdots & \alpha_k\end{pmatrix}$ where each α_i is a 2^ℓ-dimensional column vector. Let $i_1, ..., i_r \in [k]$ be column indices of M for which $M^{i_1 \cdots i_r} := \begin{pmatrix}\alpha_{i_1} & \alpha_{i_2} & \cdots & \alpha_{i_r}\end{pmatrix}$ is of full rank. Define *sub-signature* $R^{i_1 \cdots i_r}$ to consist of entries $(R_{j_1 \cdots j_n})$ of R ranging over all $j_1, ..., j_n \in \{i_1, ..., i_r\} \subset [k]$. We can define the sub-signature $G^{i_1 \cdots i_r}$ of a generator analogously. Equivalently,

$$R^{i_1 \cdots i_r} = \underline{R}(M^{i_1 \cdots i_r})^{\otimes n} \tag{15}$$

Lemma 6.2. *For a recognizer R realizable on basis M, if there exists t for which* $\operatorname{rank}(R(t)) \ge r$, *then* $\operatorname{rank}(R^{i_1 \cdots i_r}(t)) = r$.

Proof. By Lemma 2.14, $R(t) = (M^T)^{\otimes(n-1)} \underline{R}(t) M$, so $\operatorname{rank}(\underline{R}(t)) \ge r$. By (15) and Lemma 2.13, $R^{i_1 \cdots i_r}(t) = ((M^{i_1 \cdots i_r})^T)^{\otimes(n-1)} \underline{R}(t) M^{i_1 \cdots i_r}$, so $\operatorname{rank}(R^{i_1 \cdots i_r}(t)) = r$. □

For such a recognizer R, define for each $w \in [k]$ a nk^{n-1}-dimensional column vector b_w by

$$b_w = \begin{pmatrix} R_{w1\cdots11} & \cdots & R_{wk\cdots kk} & R_{1w\cdots11} & \cdots & R_{kw\cdots kk} & \cdots & R_{11\cdots1w} & \cdots & R_{kk\cdots kw}\end{pmatrix}^T \tag{16}$$

and define $A_{i_1 \cdots i_r}$ to be the $nk^{(n-1)} \times r$ matrix whose jth column is b_{i_j}.

Observation 3. $\operatorname{rank}(A_{i_1 \cdots i_r}) = r$.

Proof. $R^{i_1 \cdots i_r}(t)$ is a submatrix of $A_{i_1 \cdots i_r}$ and already has rank r by Lemma 6.2. □

Observation 4. *As M has rank r, every column α_w can be expressed as a linear combination* $\sum_{j=1}^r x_w^{ij} \alpha_{i_j}$.

Denote the $r \times k$ matrix $\left(x_w^{ij}\right)$ of these coefficients by $X_{i_1 \cdots i_r}$.

Lemma 6.3. *For each $w \in [k]$, $A_{i_1 \cdots i_r} X = b_w$ has the unique solution $X = \begin{pmatrix} x_w^{i_1} & \cdots & x_w^{i_r} \end{pmatrix}^T$.*

Proof. A solution for X exists and is unique because $\text{rank}(A_{i_1 \cdots i_r}) = r$ by Observation 3. To check that the purported solution for X is correct, pick any entry $R_{j_1 \cdots j_{t-1} w j_t \cdots j_n}$ of b_w. By definition of recognizer signatures,

$$\begin{aligned} R_{j_1 \cdots j_{t-1} w j_t \cdots j_n} &= \langle \underline{R}, \alpha_{j_1} \otimes \cdots \otimes \alpha_{j_{t-1}} \otimes \alpha_w \otimes \alpha_{j_{t+1}} \otimes \cdots \otimes \alpha_{j_n} \rangle \\ &= \langle \underline{R}, \alpha_{j_1} \otimes \cdots \otimes \alpha_{j_{t-1}} \otimes \left(\sum_{j=1}^{r} x_w^{ij} \alpha_{i_j} \right) \otimes \alpha_{j_{t+1}} \otimes \cdots \otimes \alpha_{j_n} \rangle \\ &= \sum_{j=1}^{r} x_w^{ij} \cdot \left\langle \underline{R}, \alpha_{j_1} \otimes \cdots \otimes \alpha_{j_{t-1}} \otimes \alpha_{i_j} \otimes \alpha_{j_{t+1}} \otimes \cdots \otimes \alpha_{j_n} \right\rangle \\ &= \sum_{j=1}^{r} x_w^{ij} R_{j_1 \cdots j_{t-1} i_j j_{t+1} \cdots j_n}. \end{aligned}$$

Here $\langle \cdot, \cdot \rangle$ denotes the inner product. □

The content of Lemma 6.3 is that to any such R we can get a matrix $X_{i_1 \cdots i_r}$ without needing to know the actual basis M realizing R.

Lemma 6.4. *If $\text{rank}(R_1(t)) \geq r$ for some t, then recognizers $R_1, ..., R_a$ are simultaneously realizable on some basis of rank r iff the following conditions hold:*

1. *$\text{rank}(R_1^{i_1 \cdots i_r}(t)) = r$ for some $i_1, ..., i_r \in [k]$.*
2. *There exists a unique $r \times k$ matrix $X_{i_1 \cdots i_r} = (x_w^{ij})$ such that $A_{i_1 \cdots i_r} X = b_w$ has the solution $X = \begin{pmatrix} x_w^{i_1} & \cdots & x_w^{i_r} \end{pmatrix}^T$ for each $w \in [k]$.*
3. *There exists a $2^\ell \times r$ basis $M_{(r)}$ such that the $R_j^{i_1 \cdots i_r}$ are simultaneously realizable on $M_{(r)}$ for all $j \in [a]$.*
4. *$R_j = R_j^{i_1 \cdots i_r} X_{i_1 \cdots i_r}^{\otimes n}$ for all $j \in [a]$.*

Proof. Suppose $R_1, ..., R_a$ are simultaneously realizable are some basis M. Conditions 1 and 2 follow from Lemma 6.2 and Lemma 6.3 respectively. Take $M_{(r)}$ to be $M^{i_1 \cdots i_r}$, and condition 3 follows from the definition of sub-signature. By Observation 4, $X_{i_1 \cdots i_r}$ satisfies $M_{(r)} X_{i_1 \cdots i_r} = M$, so $R_j = \underline{R}_j M^{\otimes n} = \underline{R}_j M_{(r)}^{\otimes n} X_{i_1 \cdots i_r}^{\otimes n} = R_j^{i_1 \cdots i_r} X_{i_1 \cdots i_r}^{\otimes n}$, and condition 4 follows.

Conversely, suppose conditions 1-4 hold. Condition 3 tells us that there is some $M_{(r)}$ for which $R_j^{i_1 \cdots i_r} = \underline{R}_j \left(M_{(r)}\right)^{\otimes n}$ for all $j \in [a]$. Then

$$R_j = R_j^{i_1 \cdots i_r} X_{i_1 \cdots i_r}^{\otimes n} = \underline{R}_j (M_{(r)})^{\otimes n} X_{i_1 \cdots i_r}^{\otimes n} = \underline{R}_j (M_{(r)} X_{i_1 \cdots i_r})^{\otimes n},$$

so $R_1, ..., R_a$ are simultaneously realizable on $M := M_{(r)} X_{i_1, ..., i_r}$. □

Theorem 6.5. *If recognizer signatures $R_1, ..., R_a$ and generator signatures $G_1, ..., G_b$ in matchgrid Ω are simultaneously realizable on a basis of rank r and there exists t for which $\operatorname{rank}(R_1(t)) \geq r$, then there exist recognizer signatures $\check{R}_1, ..., \check{R}_a$ and generator signatures $\check{G}_1, ..., \check{G}_b$ in matchgrid Ω' over domain size r that are simultaneously realizable on a $2^\ell \times r$ basis $M_{(r)}$. Furthermore,*

$$\operatorname{Holant}(\Omega) = \operatorname{Holant}(\Omega'). \tag{17}$$

Proof. We first construct $\check{R}_1, ..., \check{R}_a, \check{G}_1, ..., \check{G}_b$. $X_{i_1,...,i_r}$ obtained from R_1 via Lemma 6.3 has rank r, so let $X'_{i_1 \cdots i_r}$ be the $k \times k$ invertible matrix for which $X_{i_1 \cdots i_r} X'_{i_1 \cdots i_r} = \left(I_r \mid \mathbf{0}_{r \times (k-r)}\right)$, where I_r is the $r \times r$ identity matrix and $\mathbf{0}_{r \times (k-r)}$ denotes the $r \times (k-r)$ matrix consisting solely of zeroes. For each $j \in [a]$, let $R'_j = R_j(X'_{i_1,...,i_r})$, and let $\check{R}_j$ be the sub-signature $(R'_j)^{1 \cdots r}$. Likewise, for each $j \in [b]$, let G'_j be defined by $G_j = (X'_{i_1 \cdots i_r})^{\otimes n} G'_j$, and let $\check{G}_j$ be the sub-signature $(G'_j)^{1 \cdots r}$.

Claim 6.6. *For all j, $\check{R}_j = R_j^{i_1 \cdots i_r}$ and $\check{G}_j = G_j^{i_1 \cdots i_r}$.*

Proof. We need to check that

$$\check{R}_j = \underline{R}_j (M^{i_1 \cdots i_r})^{\otimes n} \tag{18}$$

$$(M^{i_1 \cdots i_r})^{\otimes n} \check{G}_j = \underline{G}_j. \tag{19}$$

Indeed,

$$\begin{aligned} R'_j &= \underline{R}_j M^{\otimes n} (X'_{i_1 \cdots i_r})^{\otimes n} \\ &= \underline{R}_j (M^{i_1 \cdots i_r})^{\otimes n} X^{\otimes n}_{i_1 \cdots i_r} (X'_{i_1 \cdots i_r})^{\otimes n} \\ &= \underline{R}_j \left(M^{i_1 \cdots i_r} \mid \mathbf{0}_{r \times (k-r)}\right)^{\otimes n}, \end{aligned}$$

proving (18). Similarly,

$$\begin{aligned} \underline{G}_j &= M^{\otimes n} G_j \\ &= M^{\otimes n} (X'_{i_1 \cdots i_r})^{\otimes n} G'_j \\ &= \left(M^{i_1 \cdots i_r} \mid \mathbf{0}_{r \times (k-r)}\right)^{\otimes n} G'_j, \end{aligned}$$

proving (19). □

We conclude that $\check{R}_1, ..., \check{R}_a, \check{G}_1, ..., \check{G}_b$ are simultaneously realizable on the basis $M_{(r)} := M^{i_1 \cdots i_r}$.

To check that the Holants agree, first note that if $R'_1, ..., R'_a, G'_1, ..., G'_b$ lie in a corresponding matchgrid Ω'', $\operatorname{Holant}(\Omega) = \operatorname{Holant}(\Omega'')$ because we're just applying a basis change from M to $MX'_{i_1 \cdots i_r}$. And $\operatorname{Holant}(\Omega') = \operatorname{Holant}(\Omega'')$ because the operation of taking sub-signatures does not lose any information in this case, i.e. $(R'_j)_\sigma = 0$ for all $\sigma \in [k]^n \backslash [r]^n$. □

For the next two results, suppose recognizer signatures $R_1, ..., R_a$ and generator signatures $G_1, ..., G_b$ in matchgrid Ω are simultaneously realizable on a basis of rank r and there exists t for which $\operatorname{rank}(R_1(t)) \geq r$.

Theorem 6.7. *If the recognizer signatures $\check{R}_1, ..., \check{R}_a$ and generator signatures $\check{G}_1, ..., \check{G}_b$ constructed in Theorem 6.5 are also simultaneously realizable on a $2^{\ell'} \times r$ basis $M'_{(r)}$ of rank r, then recognizer signatures $R_1, ..., R_a$ and generator signatures $G_1, ..., G_r$ are simultaneously realizable on the $2^{\ell'} \times k$ basis $M'_{(r)} X_{i_1 \cdots i_r}$, where $X_{i_1 \cdots i_r}$ is obtained from R_1 by Lemma 6.3.*

Proof.

$$R_j = R_j^{i_1 \cdots i_r} X_{i_1 \cdots i_r}^{\otimes n} = \check{R}_j X_{i_1 \cdots i_r}^{\otimes n} = \underline{R}(M'_{(r)} X_{i_1 \cdots i_r})^{\otimes n},$$

where the first equality holds by condition 4 of Lemma 6.4, the second by Claim 6.6, the third by definition of $M'_{(r)}$. Likewise, because

$$\check{G}_j = \left(I_r \mid \mathbf{0}_{r \times (k-r)}\right) G'_j = X_{i_1 \cdots i_r} X'_{i_1 \cdots i_r} G'_j,$$

we have that

$$\underline{G}_j = M_{(r)}^{\otimes n} \check{G}_j = (M_{(r)} X_{i_1 \cdots i_r})^{\otimes n} X'^{\otimes n}_{i_1 \cdots i_r} G'_j = (M_{(r)} X_{i_1 \cdots i_r})^{\otimes n} G_j,$$

so we conclude that $R_1, ..., R_a, G_1, ..., G_b$ are indeed simultaneously realizable on $M_{(r)} X_{i_1 \cdots i_r}$. □

We are now ready to prove Theorem 6.1.

Proof of Theorem 6.1. By Theorem 6.5, signatures $\check{R}_1, ..., \check{R}_a$ and $\check{G}_1, ..., \check{G}_b$ on domain size r are simultaneously realizable on a $2^\ell \times r$ basis.

By definition, $\check{R}_1 = R_1^{i_1 \cdots i_r}$, and because R_1 was assumed to be full-rank, Lemma 6.2 tells us that $\check{R}_1$ is full-rank. Then by the hypothesis that Theorem 1.1 has already been proven for domain size r, there exists a $2^{\lfloor \log_2 r \rfloor} \times r$ basis $M'_{(r)}$ on which $\check{R}_1, ..., \check{R}_a$ and $\check{G}_1, ..., \check{G}_b$ are simultaneously realizable. By Theorem 6.7, $R_1, ..., R_a$ and $G_1, ..., G_b$ are simultaneously realizable on $2^{\lfloor \log_2 r \rfloor} \times k$ basis $M' := M'_{(r)} X_{i_1 \cdots i_r}$. □

By Corollary 4.4 and Theorem 6.1, it remains to prove collapse theorems for holographic algorithms on domain sizes $k = 2^K$ and over bases of full rank, after which we get the following corollary.

Corollary 6.8. *Any holographic algorithm on a basis of size ℓ and domain size k not a power of 2 which uses at least one generator signature of full rank can be simulated on a basis of size at most $2^{\lfloor \log_2 k \rfloor}$.*

7 Collapse Theorem For Domain Size 2^K

The following is a direct generalization of the argument from Section 5.3 of [5], but we include it for completeness. We will take G to be a generator signature of full rank on domain size $k = 2^K$, basis M to be a $2^\ell \times 2^K$ matrix of rank 2^K, and $\underline{G} = M^{\otimes n} G$ to be the corresponding standard signature of arity $n\ell$. By Theorem 4.2 applied to the transpose of $\underline{G}(t)$, there exists a cluster $Z = s + \{e_{p_1}, ..., e_{p_K}\}$ of rows of full rank in $\underline{G}(t)$. Denote by M^Z the submatrix of M consisting of rows with indices in Z.

Lemma 7.1. *M^Z is invertible.*

Proof. The $(k, n\ell)$ cluster submatrix of $\underline{G}(t)$ of full rank whose existence is guaranteed by Theorem 4.2 is a submatrix of $M^Z G(t) (M^T)^{\otimes (n-1)}$, so M^Z has rank at least 2^K. But M^Z is a $2^K \times 2^K$ matrix, so M^Z is invertible. □

Following the notation of [5], now denote the column vector $(M^Z)^{\otimes n} G$ of dimension 2^{Kn} by $\underline{G}^{* \leftarrow Z}$ and the column vector $(M^Z)^{\otimes (t-1)} \otimes M \otimes (M^Z)^{\otimes (n-t)} \cdot G$ of dimension $2^{Kn+\ell-K}$ by $\underline{G}^{t^c \leftarrow Z}$. Because M^Z and $G(t)$ both have rank 2^K, $\underline{G}^{* \leftarrow Z}$ and $\underline{G}^{t^c \leftarrow Z}$ also have rank 2^K. We check that these can be realized as standard signatures.

Lemma 7.2. *$\underline{G}^{* \leftarrow Z}$ is the standard signature of a generator matchgate of arity Kn.*

Proof. Take the matchgate G, and in each block, attach an edge of weight 1 to external node i ($1 \leq i \leq \ell$) if $s_i = 1$. In the matchgate G' we get from these operations, designate external nodes $p_1, ..., p_K$ in each block as the new external nodes of G'. The resulting matchgate realizes $\underline{G}^{*\leftarrow Z}$. □

Lemma 7.3. *$\underline{G}^{t^c\leftarrow Z}$ is the standard signature of a generator matchgate of arity $Kn - K + \ell$.*

Proof. The proof of Lemma 7.3 is almost identical to that of Lemma 7.2, except block t is treated differently. Take the matchgate G, and in each block except the t-th one, attach an edge of weight 1 to external node i ($1 \leq i \leq \ell$) if $s_i = 1$. In the matchgate G' we get from these operations, take the external nodes to be all ℓ external nodes in block t, as well as nodes $p_1, ..., p_K$ in every other block. The resulting matchgate realizes $\underline{G}^{t^c\leftarrow Z}$. □

Now define $T = M(M^Z)^{-1}$. Here is the key step of the collapse theorem, making use of Theorem 5.1.

Lemma 7.4. *T is the standard signature of a K-input, ℓ-output transducer.*

Proof. We first express T in terms of $\underline{G}^{*\leftarrow Z}$ and $\underline{G}^{t^c\leftarrow Z}$. If the entries of $\underline{G}^{t^c\leftarrow Z}$ are indexed by $(i_{1,1}\cdots i_{1,K})\cdots(i_{t-1,1}\cdots i_{t-1,K})(i'_1\cdots i'_\ell)(i_{t+1,1}\cdots i_{t+1,\ell})\cdots(i_{n,1}\cdots i_{n,K})$, denote by $\underline{G}^{t^c\leftarrow Z}(t)$ the matrix form of $\underline{G}^{t^c\leftarrow Z}$ in which the rows are indexed by $i'_1\cdots i'_\ell$ and the columns are indexed by $(i_{1,1}\cdots i_{1,K})\cdots(i_{t-1,1}\cdots i_{t-1,K})(i_{t+1,1}\cdots i_{t+1,\ell})\cdots(i_{n,1}\cdots i_{n,K})$.

Observe that

$$\underline{G} = M^{\otimes n}G = T^{\otimes n}(M^Z)^{\otimes n}G = T^{\otimes n}\underline{G}^{*\leftarrow Z}$$

so that

$$\underline{G}^{t^c\leftarrow Z} = (T^Z)^{\otimes(t-1)} \otimes T \otimes (T^Z)^{\otimes(n-t)}\underline{G}^{*\leftarrow Z}. \tag{20}$$

Putting both sides of (20) in matrix form, we conclude that

$$\underline{G}^{t^c\leftarrow Z}(t) = T\underline{G}^{*\leftarrow Z}(t). \tag{21}$$

Applying Theorem 5.1 to the arity-Kn standard signature $\underline{G}^{*\leftarrow Z}$, we have a recognizer whose standard signature $\underline{R}$ satisfies $\underline{G}^{*\leftarrow Z}(t)\underline{R}(t) = I_{2^K}$. Right-multiplying both sides of (21) by $\underline{R}(t)$, we find that

$$\underline{G}^{t^c\leftarrow Z}(t)\underline{R}(t) = T.$$

Say that the generator realizing $\underline{G}^{t^c\leftarrow Z}$ as a standard signature has external nodes $X_{i,1}$, $X_{i,2}$, ..., $X_{i,K}$ in block i for each $i \neq t$, and external nodes $Y_{t,1}$,..., $Y_{t,\ell}$ in block t. Say that the generator realizing $\underline{R}$ as a standard signature has external nodes $Z_{i,1}, ..., Z_{i,K}$ in each block i.

Construct the transducer Γ realizing T as a standard signature by connecting $X_{i,j}$ with $Z_{i,j}$ for all $i \neq t$, $j \in [K]$. Designate $Y_{t,1}, ..., Y_{t,\ell}$ to be the output nodes of Γ and $Z_{t,1}$, ..., $Z_{t,K}$ to be the input nodes of Γ. □

From Theorem 7.4 we obtain the collapse theorem for domain size 2^K.

Theorem 7.5. *Any holographic algorithm on a basis of size ℓ and domain size 2^K which uses at least one generator signature of full rank can be simulated on a basis of size K.*

Proof. Suppose the holographic algorithm in question uses signatures R_i, G_j ($1 \leq i \leq r$, $1 \leq j \leq g$) defined by $\underline{R_i}M^{\otimes m_i} = R_i$ and $\underline{G_j} = M^{\otimes n_j}G_j$ over basis M. Say that G_1 has full rank, and let $Z = s + \{e_{p_1}, ..., e_{p_K}\}$ denote the full-rank (K, ℓ)-cluster of rows in G_1 which must exist by Theorem 4.3. By Lemma 7.4, $T := M(M^Z)^{-1}$ is the standard signature of some transducer matchgate Γ. Let $\underline{R'_i} = \underline{R_i}T^{\otimes m_i}$ and $\underline{G'_j} = \underline{G_j}^{*\leftarrow Z}$; by Lemma 2.4, $\underline{R'_i}$ is the standard signature of some recognizer, and by Lemma 7.3, $\underline{G'_j}$ is the standard signature of some generator. We conclude that the R_i, G_j can be simultaneously realized on the basis M^Z of size K. □

8 Acknowledgments

I would like to thank Professor Leslie Valiant for patiently advising me throughout the course of this project and providing much valuable feedback as well as suggestions on papers to read. I am also extremely indebted to Professor Jin-Yi Cai for his incredibly extensive and important comments on drafts of this paper.

References

[1] J.-Y. Cai and V. Choudhary. Some results on matchgates and holographic algorithms, *Automata, languages, and programming. Part 1, Lecture Notes in Comput. Sci.*, vol 4051, Springer, Berlin, 2006, pp. 703-714.

[2] J.-Y. Cai and V. Choudhary. Valiant's holant theorem and matchgate tensors, *Theory and applications of models of computation, Lecture Notes in Comput. Sci.*, vol. 3959, Springer, Berlin, 2006, pp. 248-261.

[3] J.-Y. Cai and V. Choudhary. Valiant's holant theorem and matchgate tensors, *Theoret. Comput. Sci.* **384** (2007), no. 1, 22-32.

[4] J.-Y. Cai, V. Choudhary, and P. Lu. On the theory of matchgate computations, *Theory Comput. Syst.*, 45(1):108132, 2009. Preliminary version in CCC'07.

[5] J.-Y. Cai and Z. Fu. (2014). A collapse theorem for holographic algorithms with matchgates on domain size at most 4. *Information and Computation*, 239, 149-169.

[6] J.-Y. Cai and A. Gorenstein. Matchgates revisited. *CoRR*, abs/1303.6729, 2013.

[7] J.-Y. Cai and P. Lu. Holographic algorithms: from art to science, *STOC '07– Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, AMC, New York, 2007, pp. 401-410.

[8] J.-Y. Cai and P. Lu: Holographic algorithms: the power of dimensionality resolved. In: Arge, L., Cachin, C., Jurdzinski, T. et al. (eds.) *Proceedings of ICALP*. Lecture Notes in Computer Science, vol. 4596, pp. 631642. Springer, Berlin (2007)

[9] J.-Y. Cai and P. Lu. On symmetric signatures in holographic algorithms, *STACS 2007, Lecture Notes in Comput. Sci.*, vol. 4393, Springer, Berlin, 2007, pp. 429-440.

[10] J.-Y. Cai and P. Lu. Basis collapse in holographic algorithms, *Computational Complexity* **17** (2008), no. 2, 254-281.

[11] J.-Y. Cai, P. Lu, and M. Xia. Holographic reduction, interpolation and hardness. *Computational Complexity*, 21(4):573604, 2012.

[12] Z. Fu, F. Yang: Holographic algorithms on bases of rank 2. `http://arxiv/abs/1303.7361`. In submission.

[13] P.W. Kasteleyn: The statistics of dimers on a lattice. Physica 27, 12091225 (1961)

[14] J. Landsberg, J. Morton, and S. Norine. Holographic algorithms without matchgates. preprint arXiv:0904.0471, 2009.

[15] A. Li and M. Xia: A theory for Valiant's matchcircuits (extended abstract). *In Proc. 25th Symp. Theoretical Aspects of Comp. Sci. (STACS'08)*, pp. 491502. Schloss Dagstuhl, 2008.

[16] K. Mulmuley and M. Sohoni. Geometric complexity theory I: An approach to the P vs. NP and related problems. In *SIAM J. Comput.*, vol 31, no. 2, 496-526, 2001.

[17] H.N.V. Temperley and M.E. Fisher. Dimer problem in statistical mechanicsan exact result. Philos. Mag. 6, 10611063 (1961)

[18] S.P. Vadhan. The complexity of counting in sparse, regular and planar graphs, *SIAM J. Comput.* 31(2): 398427, 2001.

[19] L.G. Valiant. Expressiveness of matchgates, *Theoret. Comput Sci.* **289** (2002), no. 1, 457-471.

[20] L.G. Valiant. Quantum circuits that can be simulated classically in polynomial time. SIAM J. Comput., 31(4):12291254, 2002. Preliminary version in STOC'01.

[21] L.G. Valiant. Accidental algorithms. In *FOCS*, pages 509517. IEEE Computer Society, 2006.

[22] L.G. Valiant. Holographic algorithms (extended abstract). In *Proc. 45th IEEE Symposium on Foundations of Computer Science*, pp. 306315 (2004).

[23] L.G. Valiant. Holographic circuits, *Automata, languages and programming, Lecture Notes in Comput. Sci.*, vol. 3580, Springer, Berlin, 2005, pp. 1-15.

[24] L.G. Valiant. Some observations on holographic algorithms, Proc. 9th Latin American Theoretical Informatics Symposium, LATIN 2010. LNCS, Vol 6034 Springer-Verlag (2010), 577-590.

[25] M. Xia, P. Zhang, and W. Zhao, The complexity of counting on 3-regular planar graphs, *Theoret. Comput. Sci.*, 384 (2007), pp. 111125.

A Planarizing Matchgates

In the proof of Lemma 5.3, we made several initial constructions of transducers to achieve certain row and column operations but noted that those constructions, specifically those shown in Figures 2b, 2c, 2d, needed to be modified because they were not planar. Following the technique of Cai and Gorenstein [6], we planarize those matchgates by replacing every edge crossing with the so-called *crossover gadget* X shown in Figure 3.

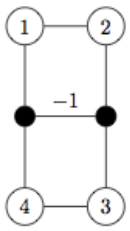

Figure 3: Crossover gadget. Unlabeled edges are of weight 1; labeled vertices are external nodes. Figure from [6].

Because the standard signature $\underline{X}$ of the crossover gadget is given by $\underline{X}^{0000} = 1$, $\underline{X}^{0101} = 1$, $\underline{X}^{1010} = 1$, $\underline{X}^{1111} = -1$, and $\underline{X}^{\sigma} = 0$ for all other $\sigma \in \{0,1\}^4$ so that the standard signature

remains invariant under any cyclic permutation of the external nodes, the orientation of the copy of X placed over an edge crossing does not matter.

We first make precise our operation of planarizing matchgates, following the terminology of [6]. If an edge $\{u, v\}$ of weight w crosses t other edges, replace each of the t crossings by a crossover gadget and replace the edge by $t + 1$ edges connecting adjacent crossover gadgets. Of these $t + 1$ edges, assign t of them to have weight 1 and the remaining one to have weight w. Call the union of the $t + 1$ edges the *u-v-passage*.

Given a non-planar matchgate Γ, denote the matchgate obtained from planarizing Γ by Γ'.

Observation 5. *Let M be a perfect matching of Γ' whose contribution c to* PerfMatch(Γ') *is nonzero, and let $M' \subset M$ denote the edges not belonging to crossover gadgets. Then M' is the union of u-v-passages corresponds to a perfect matching of Γ whose contribution to* PerfMatch(Γ) *is $\pm c$.*

Proof. If an edge incident to one of the external nodes, say node 1, of a crossover gadget is present in M, then the edge incident to node 3 of the crossover gadget must be present in M as well, as $\underline{X}^{\sigma} = 0$ if $\sigma_1 \neq \sigma_3$. We conclude that M is a union of u-v-passages. The corresponding perfect matching of Γ has contribution $\pm c$ because each of the nonzero entries of $\underline{X}$ is ± 1. □

We need to verify that the entries of $\underline{\Gamma}$ and $\underline{\Gamma'}$ are equal except for a select number of entries which differ by a factor of -1.

Lemma A.1. *Let Γ be the K-input, K-output transducer shown in Figure 2b with signature $L_3^{j,k}$. There exists a planar matchgate whose standard signature agrees with $\underline{\Gamma}$ on the main diagonal entries and entry $(1^K \oplus e_j \oplus e_k, 1^K)$, and agrees with $\underline{\Gamma}$ everywhere else up to sign.*

Proof. Take the desired matchgate to be Γ'. Note that every subgraph of Γ has at most one perfect matching. In other words, each entry of $\underline{\Gamma}$ arises from at most a single perfect matching. Therefore, by Observation 5, $\underline{\Gamma}$ and $\underline{\Gamma'}$ agree everywhere up to sign. Now consider any main diagonal entry $\underline{\Gamma'}^{\sigma}_{\sigma} =$ PerfMatch($\Gamma' \backslash Z$). If M is a perfect matching of $\Gamma' \backslash Z$ making a nonzero contribution to PerfMatch($\Gamma' \backslash Z$), it corresponds to a perfect matching of $\Gamma \backslash Z$ making a nonzero contribution to PerfMatch($\Gamma \backslash Z$). But the only such perfect matching does not contain the edge between left node j to and left node k. Thus, the contribution of this matching and that of M are both equal to 1.

If $\underline{\Gamma}$ and $\underline{\Gamma'}$ disagree on entry $(1^K \oplus e_j \oplus e_k, 1^K)$, modify Γ by multiplying the weight of the edge connecting left nodes j and k by -1, and take the desired matchgate to be the corresponding Γ'. □

Lemma A.2. *Let Γ be the K-input, K-output transducer shown in Figure 2c with signature L_4. There exists a planar matchgate whose standard signature agrees with $\underline{\Gamma}$ everywhere up to sign.*

Proof. Take the desired matchgate to be Γ'. As in the proof of Lemma A.1, every subgraph of Γ has at most one perfect matching, so we already know $\underline{\Gamma}$ and $\underline{\Gamma'}$ agree everywhere up to sign. □

Lemma A.3. *Let Γ be the K-input, K-output transducer shown in Figure 2d with signature L_5^j. There exists a planar matchgate whose standard signature agrees with $\underline{\Gamma}$ on the main diagonal entries and entry $(1^K \oplus e_j, 1^K \oplus e_{q_{i+1}})$, and agrees with $\underline{\Gamma}$ everywhere else up to sign.*

Proof. Take the desired matchgate to be Γ'. As in the proof of Lemma A.1, every subgraph of Γ has at most one perfect matching and $\underline{\Gamma}$ and $\underline{\Gamma'}$ agree on the main diagonal entries. If they disagree on entry $(1^K \oplus e_j, 1^K \oplus e_{q_{i+1}})$, modify Γ by multiplying the weight of the edge connecting left node j to right node $i + 1$ by -1, and take the desired matchgate to be the corresponding Γ'. □